

UNIVERSIDADE FEDERAL DE MINAS GERAIS
Instituto de Ciências Exatas
Programa de Pós-Graduação em Ciência da Computação

Alexandre Magno de Sousa

Analyzing and Modeling User Curiosity in Online Information Services

Belo Horizonte
2024

Alexandre Magno de Sousa

Analyzing and Modeling User Curiosity in Online Information Services

Final Version

Dissertation presented to the Graduate Program in Computer Science of the Federal University of Minas Gerais in partial fulfillment of the requirements for the degree of Doctor in Computer Science.

Advisor: Jussara Marques de Almeida

Co-Advisor: Flavio Vinicius Diniz de Figueiredo

Belo Horizonte
2024

Sousa, Alexandre Magno de.

S725a Analyzing and modeling user curiosity in online information
Services [recurso eletrônico] / Alexandre Magno de Sousa.–
2024.

1 recurso online (192 f. il, color.) : pdf.

Orientadora: Jussara Marques de Almeida.

Coorientador: Flávio Vinicius Diniz de Figueiredo.

Tese (Doutorado) - Universidade Federal de Minas
Gerais, Instituto de Ciências Exatas, Departamento de
Ciências da Computação.

Referências: f. 165-182

1. Computação – Teses. 2. Redes sociais on-line – Teses.
3. Serviços de informação – Estudo de usuários - Teses.
4. Teoria da informação em psicologia – Webometria – Teses.
I. Almeida, Jussara Marques de. II. Figueiredo, Flávio Vinicius
Diniz de. III. Universidade Federal de Minas Gerais, Instituto de
Ciências Exatas, Departamento de Computação. IV. Título.

CDU 519.6*22(043)



UNIVERSIDADE FEDERAL DE MINAS GERAIS
INSTITUTO DE CIÊNCIAS EXATAS
PROGRAMA DE PÓS-GRADUAÇÃO EM CIÊNCIA DA COMPUTAÇÃO

FOLHA DE APROVAÇÃO

Analyzing and Modeling User Curiosity in Online Information Services

ALEXANDRE MAGNO DE SOUSA

Tese defendida e aprovada pela banca examinadora constituída pelos Senhores:

Documento assinado digitalmente



JUSSARA MARQUES DE ALMEIDA GONCALVES

Data: 18/10/2024 07:25:19-0300

Verifique em <https://validar.iti.gov.br>

PROFA. JUSSARA MARQUES DE ALMEIDA GONÇALVES - Orientadora
Departamento de Ciência da Computação - UFMG

PROF. FLAVIO VINICIUS DINIZ DE FIGUEIREDO - Coorientador
Departamento de Ciência da Computação - UFMG

Documento assinado digitalmente



ANA PAULA COUTO DA SILVA

Data: 18/10/2024 11:20:33-0300

Verifique em <https://validar.iti.gov.br>

PROFA. ANA PAULA COUTO DA SILVA
Departamento de Ciência da Computação - UFMG

Documento assinado digitalmente



MARIO SERGIO FERREIRA ALVIM JUNIOR

Data: 20/10/2024 15:10:53-0300

Verifique em <https://validar.iti.gov.br>

PROF. MARIO SÉRGIO FERREIRA ALVIM JÚNIOR
Departamento de Ciência da Computação - UFMG

Documento assinado digitalmente



THIAGO HENRIQUE SILVA

Data: 18/10/2024 09:39:09-0300

Verifique em <https://validar.iti.gov.br>

PROF. THIAGO HENRIQUE SILVA
Departamento Acadêmico de Informática - UTFPR

Documento assinado digitalmente



ALEX BORGES VIEIRA

Data: 18/10/2024 09:49:26-0300

Verifique em <https://validar.iti.gov.br>

PROF. ALEX BORGES VIEIRA
Departamento de Ciência da Computação - UFJF

Belo Horizonte, 26 de junho de 2024.

To my wife Fernanda, who celebrated my realizations, supported me in my failures and patiently endured the long nights of research. This achievement would not have been possible without your understanding.

Acknowledgments

First of all, I thank God for giving me strength and supporting me to complete this PhD, which has been a journey of patience and humility. Next, I would like to thank my advisors Jussara and Flávio for their important and valuable advice, especially Professor Jussara for her guidance, understanding and patience during the development of this thesis. I would also like to thank the Universidade Federal de Minas Gerais, represented by the faculty and administrative staff who carry out the task of professional administration of the postgraduate program. In particular, I would like to thank Professor Newton José Vieira for encouraging me to start my PhD when he literally led me to Professor Jussara's office.

I cannot fail to mention my colleagues in the research lab who supported me and were part of my growth during my PhD: Carlos, Fabiano, Glauber and Júlio, to whom my thanks are due. I would also like to thank my colleagues at the Social Computing Laboratory (LoCuS) and the Universidade Federal de Ouro Preto.

I would also like to thank my parents, José Eustáquio and Rosimeire, for their intercession and encouragement during this journey. And finally, and most importantly, I would like to thank my wife Fernanda for her affection, understanding, companionship and help throughout my doctoral journey, as well as in the difficult moments that surprised us along the way.

“Not only that, but we also rejoice in our sufferings, because we know that suffering produces perseverance; perseverance, character; and character, hope.”
(Romans 5:3–4)

Resumo

Na atual era da economia da informação, a curiosidade, componente fundamental do comportamento humano, exerce um papel de extrema importância, uma vez que diversos serviços de informação (e.g., busca, recomendação) são projetados para atrair e manter a atenção do usuário. Recentemente, muitos esforços têm surgido para introduzir modelos de curiosidade em serviços de informação online visando melhorar a personalização de seus resultados e, em última instância, a satisfação de seus usuários. Um modelo de particular interesse, fundamentado em teorias da Psicologia, é a curva de Wundt, que expressa como a curiosidade de uma pessoa varia em função do estímulo ao qual ela é exposta. O estímulo, por sua vez, é definido como uma combinação de variáveis colativas, tais como novidade, complexidade, conflito e incerteza, que capturam fatores que governam como a curiosidade humana é estimulada.

O objetivo desta tese é investigar a hipótese de que modelos de curiosidade computacional podem ser utilizados para obter uma compreensão mais fundamental do processo de disseminação de informação online e para projetar serviços de informação personalizados mais eficazes. Para tanto, propõe-se explorar múltiplas variáveis colativas que compõem o estímulo de curiosidade, um esforço ainda pouco explorado na literatura, em especial no contexto de serviços de informação online. A investigação é feita a partir de dois estudos de caso relacionados à disseminação de informação. O primeiro aborda o papel da curiosidade no consumo de informação, especificamente músicas na plataforma LastFM. O segundo foca na produção de informação e aborda o papel da curiosidade no compartilhamento de conteúdo em grupos do WhatsApp.

Os resultados do primeiro estudo de caso mostram que as métricas de novidade, complexidade e incerteza capturam componentes complementares da curiosidade na maioria dos casos. Por meio dessas métricas, foram encontrados quatro perfis de curiosidade distintos em nível de acesso do consumo da informação. Também foram encontrados padrões que mostram como a curiosidade do usuário muda por meio desses diferentes perfis. Os resultados também mostram que os usuários dão preferência aos conteúdos que são diferentes do que eles têm consumido recentemente, talvez para evitar o tédio gerado por conteúdos familiares.

Além disso, para o segundo estudo de caso, argumenta-se que é necessário considerar a influência social como outro componente do estímulo de curiosidade, em especial nas plataformas de mídias sociais online. Dessa forma, são propostas novas métricas para capturar este aspecto tanto em nível individual como de grupo para mostrar que elas

capturam componentes da estimulação de curiosidade que não têm sido considerados até agora. Em nível do compartilhamento de mensagem, foram identificados três perfis e, para quase 30% dos casos, a influência social representa um importante papel na disseminação da informação. Além disso, a influência social muda muito ao longo do tempo, pois é um componente da curiosidade do usuário que responde à dinâmica das interações dentro do grupo. De fato, os resultados mostram que o impacto da curiosidade social está entre 50% e 100% para apenas 10% dos grupos, pois depende fortemente do engajamento do usuário por meio do compartilhamento de conteúdo.

Finalmente, os resultados mostram que, para uma parte significativa de usuários, a curva do perfil de curiosidade é multimodal, refletindo os padrões complexos e dinâmicos do comportamento de usuários de plataformas de mídias sociais. Esses resultados superam as deficiências de estudos anteriores que usam um formato unimodal para essa curva.

Palavras-chave: curiosidade; influência social; métricas de teoria da informação; comportamento do usuário; disseminação da informação; serviços de informação.

Abstract

In the current era of information economy, curiosity, a fundamental component of human behavior, plays a key role as various information services (e.g., search, recommendation) are designed to attract and maintain user attention. Recently, many efforts have emerged to introduce curiosity models in the design of online information services, aiming to improve result personalization and, ultimately, user satisfaction. A model of particular interest, based on theories of Psychology, is the Wundt's curve, which expresses how a person's curiosity varies depending on the stimulus to which she is exposed. Stimulus, in turn, is defined as a combination of collative variables, such as novelty, complexity, conflict and uncertainty, which capture factors governing how human curiosity is stimulated.

The goal of this dissertation is to investigate the hypothesis that computational curiosity models can be used to gain a more fundamental understanding of the online information dissemination process and to design more effective personalized information services. To that end, we propose to explore multiple collative variables as components of curiosity stimulation, an effort that has been little explored in the literature, especially in the online information services domain. The investigation is based on two case studies related to information dissemination. The first one addresses the role of curiosity in the information consumption, specifically music on the LastFM platform. The second one focuses on information production and addresses the role of curiosity as a driving force behind content sharing in WhatsApp groups.

The results in our first case study show that our metrics for novelty, complexity and uncertainty capture complementary components of curiosity in most cases. Using these metrics, we uncover four distinct profiles of curiosity at the access level of information consumption. We also uncovered dynamic patterns that show how users' curiosity changes across these different profiles. Also, the users prefer content that is different from what they have consumed in the recent past, perhaps to avoid boredom from repeated consumption.

Furthermore, for the second case study, we argue that it is necessary to consider social influence as another component of curiosity stimulus, especially on social media platforms. Therefore, we also propose new metrics to capture this aspect at both individual and group levels and show that they capture components of curiosity stimulation that have not been considered so far. At the message sharing level, we uncovered three profiles and in almost 30% of the cases, social influence plays an important role in information dissemination. Social influence changes greatly over time as it is a component of the user's

curiosity that responds to the dynamics of interaction within the group. In fact, we found that the impact of social curiosity is between 50% and 100% in only 10% of groups, as it strongly depends on user engagement through sharing content.

Finally, we found that the curiosity profile curve is multimodal for a significant proportion of users, reflecting the complex and dynamic behavioral patterns of social media platforms users. This overcomes the shortcomings of previous studies that use a unimodal curve for the curiosity profile.

Keywords: human curiosity; social influence; information theory metrics; user behavior profiles; information dissemination; information services.

List of Figures

1.1	Pictorial representation of research questions of dissertation.	23
2.1	Mind mapping of the domains which compose this dissertation.	26
2.2	Wundt's curve and its zones of curiosity stimulation.	28
2.3	Applications of computational curiosity models.	32
3.1	Overview of the problem statement.	48
3.2	Main elements of our first case study: online music consumption in LastFM. .	49
3.3	Main elements of our second case study: WhatsApp groups.	50
3.4	Methodological areas and tasks of this dissertation: (a) <i>blue blocks</i> represent the tasks and (b) <i>gray documents</i> represent the artifacts (i.e., outputs) of the respective tasks.	54
3.5	Structure for modeling stimulus of curiosity in social media with its key elements: (a) <i>user</i> as central element; (b) <i>content items</i> of information; (c) <i>social connections</i> in the system; (d) <i>user's actions</i> as different types of actions including social interactions; (e) <i>history of user accesses</i> recorded by the system.	58
4.1	Levels of content categorization for song (content items).	62
4.2	Fraction of redundant cases for each pair of metrics.	71
4.3	Determining the number of access profile.	73
4.4	Radar-chart with normalized stimulus metrics per cluster.	74
4.5	Example of UBMG and Silhouette index for number of clusters.	75
4.6	User-level curiosity stimulation profiles: state transition diagrams.	78
4.7	Mixed Wundt's curve: examples of 4 users.	80
4.8	PDF of Access Profile per Beta distribution.	81
4.9	Access Profile: examples of four users.	82
4.10	CDF of fraction of accesses per profile per user.	82
4.11	Division of the user's activity period into shorts time period (e.g., weeks). . . .	83
4.12	Example of user with 2-Beta mixture: Wundt's curve and access profile. . . .	85
4.13	Example of user with 3-Beta mixture: Wundt's curve and access profile. . . .	86
4.14	Design of Null Model 0 and 1.	87
4.15	Design of Null Model 2.	88
4.16	CDFs of the differences per access profile for <i>artistNovelty</i>	89
4.17	CDFs of the differences per access profile for <i>songNovelty</i>	90

4.18	CDFs of the differences per access profile for <i>genreNovelty</i>	91
4.19	CDFs of the differences per access profile for <i>genreUncertainty</i>	92
5.1	Contingency table $\mathbf{S}_{t,g}^-$ computed when user 3 shares message 12 in Figure 3.3 of Chapter 3.	104
5.2	Fraction of redundant (user, group) cases for each pair of metrics.	119
5.3	Diversity of social curiosity stimuli across message sharings and users.	120
5.4	Example UBMG for a user u with a sequence of message-level stimulation pattern equal to PPIDDDIPPIPIIDDDIPP in a group g (P for independent, I for indirect and D for dependent).	123
5.5	Silhouette score: determining the number of UBMG clusters.	124
5.6	Results from Principal Component Analysis applied to $k = 5, 9$ and 13 clusters.	124
5.7	User-level social curiosity stimulation profiles: state transition diagrams (UBMGs), each element is a (user, group) pair.	126
5.8	Time series of both metrics of social curiosity for different (user,group) pairs.	127
5.9	Diversity of social stimulation of selected users in different groups: the same user in different groups is represented by the same color.	128
5.10	Social curiosity stimulation across groups: group mutual information versus entropy of curious user, measured at the the time of the last message in the group (groups of user 0 in blue).	130
5.11	Group-Level Social Curiosity Stimulation: CDFs of (a) group lifespan and (b) the impact of social curiosity, measured when group reaches 10^{th} , 50^{th} and 90^{th} percentiles of user participation; (c) Time series of the impact of social influence of 3 example groups; (d)–(f) Time series of group-level metrics for the same 3 groups.	131
5.12	Graph representation of direct social influence among members of two selected groups (node diameters are proportional to out-degrees, red/blue edges refer to strong/weak social influence, thus strong/weak social curiosity stimulation.	134
5.13	Mixed Wundt’s curve: examples of 4 users.	137
5.14	CDF of fraction of social curiosity stimulation per profile per user.	138
5.15	Profile of social curiosity stimulation: examples of four users.	139
5.16	Example of user with 2-Beta mixture: Wundt’s curve and stimulation profile.	140
5.17	Example of user with 3-Beta mixture: Wundt’s curve and stimulation profile.	141
5.18	Design of Null Model 0 and 1.	143
5.19	Design of Null Model 2.	144

5.20	CDFs of social metrics for Null Models, <i>Maximum Direct Influence</i> and <i>Maximum Indirect Influence</i> , respectively. Each one of CDFs denotes the identified message-level social curiosity stimulation profiles: (a) <i>Independent</i> depicted by blue line; (b) <i>Indirect</i> represented by red line; at end, (c) <i>Dependent</i> illustrated by dashed black line.	145
5.21	CDFs of stimulus values of <i>Maximum Direct Influence</i> and <i>Maximum Indirect Influence</i> for <i>Dependent</i> Profile and <i>Indirect</i> profile: (a) Real stimulus values are depicted in <i>green line</i> ; (b) random stimulus values of Null Model 0 denoted by <i>blue line</i> ; (c) random stimulus values of Null Model 1 in <i>red line</i> ; at last, (d) random stimulus values of Null Model 2 in <i>dashed black line</i>	146
6.1	Diagram with the major research questions and supplementary research questions associated with each case study.	157
A.1	Relation of <i>songNovelty</i> , <i>artistNovelty</i> and <i>genreNovelty</i> and their respective probabilities for the axiom MONO-SURP.	185
A.2	Relation of <i>instComplex</i> , <i>overallComplex</i> and <i>genreConflict</i> and their respective probabilities for the axiom MONO-SURP.	185
A.3	Relation of <i>genreUncertainty</i> and D_{KL} value for the axiom MAX-MIN.	186
A.4	Relation of <i>userNovelty</i> and <i>catNovelty</i> and their respective probabilities for the axiom MONO-SURP.	186
A.5	Relation of <i>userUncertainty</i> and <i>catUncertainty</i> and their respective D_{KL} values for the axiom MAX-MIN.	187
A.6	Relation of <i>userConflict</i> , <i>catConflict</i> and <i>catComplex</i> and their respective probabilities for the axiom MONO-SURP.	188
B.1	The axiom MONO-PMI: relationship between the metrics of direct social influence and the ratio $P_{t,g}^{\leftarrow}(U=u S=s)/P_{t,g}^{\leftarrow}(U=u)$	191
B.2	Relationship between the metrics of indirect social influence.	192
B.3	Relationship between the metrics and the axiom BOUND-MI and BND-PTMI.	192

List of Tables

2.1	Summary of prior efforts to introduce curiosity into recommendation systems based on <i>behavioral data</i>	39
2.2	Overview of prior work and our present effort to estimate social influence. . .	44
4.1	Basic notation used to derive the curiosity stimulation metrics.	63
4.2	Metrics of curiosity stimulation for music consumption in LastFM.	69
4.3	Statistics of LFM-1B Dataset.	70
4.4	User-level curiosity stimulation profiles.	76
4.5	Validation of users.	79
4.6	Statistics of Access Profile values per Beta (mixture) distribution.	80
4.7	Results of experiments for temporal analysis with I.C. of 95%.	84
5.1	Main notation used to derive the curiosity stimulation metrics.	102
5.2	Computing the metrics of social curiosity for curious user 3 at time $t=3pm$ (reference: Figure 5.1).	110
5.3	Metrics of curiosity stimulation for content sharing in WhatsApp groups. . . .	112
5.4	Message-level social curiosity stimulation profiles.	122
5.5	User-level social curiosity stimulation profiles: percentage of messages in each UBMG state.	125
5.6	Test of statistical difference of social stimulation metrics of the same user on different groups (users 0–7 refer to selected users in Figure 5.3b).	129
5.7	Three example messages shared by the most active user in group “Science, Religion and Politics”.	133
5.8	Validation of users.	136
5.9	Statistics for profile of stimulation per Beta (mixture) distribution.	138
5.10	Results of experiments for temporal analysis with C.I. of 95%.	141

Contents

1	Introduction	18
1.1	Motivation	19
1.2	Problem Statement	21
1.3	Objectives	22
1.4	Contributions	23
1.5	Outline	25
2	Background and Related Work	26
2.1	Psychology of Curiosity	27
2.2	Computational Models of Curiosity	31
2.2.1	Curiosity Models in Online Information Systems	33
2.2.2	Principles of Curiosity Models from Recommendation Systems	34
2.2.2.1	Inferring User Curiosity from Questionnaire Data	35
2.2.2.2	Inferring User Curiosity from Behavioral Data	36
2.3	Social Influence	40
2.4	Summary	45
3	Problem Statement and Methodology	47
3.1	Case Studies	48
3.2	Assumptions	51
3.3	Problem Definition	53
3.4	Methodology	54
3.5	Fundamental structure and key elements	58
3.6	Summary	60
4	Analyzing and Modeling User Curiosity in Online Content Consumption: A LastFM Case Study	61
4.1	Introduction	61
4.2	Modeling User Curiosity	62
4.2.1	Notation	62
4.2.2	Computing Stimuli	64
4.2.3	The LastFM Dataset	70
4.3	Understanding Metrics	70
4.3.1	Variable Selection	71

4.3.2	Clustering by Access Profile	72
4.3.3	Clustering by User Behavior	74
4.4	Curiosity Models	78
4.4.1	Classes of Wundt's Curve	79
4.4.2	Curiosity and Access Patterns	80
4.4.3	Temporal Analysis	83
4.5	Validation of Curiosity Model	86
4.6	Summary	92
5	Metrics of Social Curiosity: The WhatsApp Case	95
5.1	Introduction	95
5.2	Initial Considerations	97
5.3	WhatsApp	98
5.4	Novel Metrics of User Curiosity	99
5.4.1	Notation	100
5.4.2	Social Curiosity	101
5.4.3	Other Collative Variables	112
5.5	Curiosity Stimulation in WhatsApp Groups	117
5.5.1	Relationships among Collative Variables	118
5.5.2	Diversity and Dynamics of Social Curiosity	120
5.5.2.1	Social Curiosity at the Message Level	121
5.5.2.2	Social Curiosity at the User Level	122
5.5.2.3	Social Curiosity at the Group Level	128
5.5.2.4	Some Illustrative Examples	133
5.6	Curiosity Models	136
5.6.1	Classes of Wundt's Curve	136
5.6.2	Curiosity and Stimulation Patterns	137
5.6.3	Temporal Analysis	140
5.7	Validation of Curiosity Model	142
5.8	Limitations of our Study and their Implications	147
5.8.1	Lack of Content Representation	148
5.8.2	Different Dimensions of Online Information Dissemination at the Level of User Behavior	149
5.8.3	Media Categorization	150
5.8.4	Focus on Intra-group Curiosity Stimulation	151
5.8.5	Modeling of Temporal Dynamics	151
5.9	Summary	152
6	Conclusions and Future Research Directions	156
6.1	Current Achievements	156

6.1.1	RQ1: Multiple Collative Variables and Discovery of User Curiosity Profiles	158
6.1.2	RQ2: Social Influence as Component of Human Curiosity Stimulation	159
6.1.3	RQ3: Investigating Wundt’s Curve as Model of Human Curiosity . .	160
6.2	Future Research Directions	161
6.2.1	Application of Curiosity Models in the Information Dissemination .	161
6.2.2	Application of Curiosity Models in Recommendation Systems . . .	162
6.3	Publications	164
Bibliography		165
Appendix A Axioms for Traditional Collative Variables		183
A.1	The LastFM Case Study: examples	184
A.2	The WhatsApp Case Study: examples	186
Appendix B Axioms for Social Curiosity Variables		189
B.1	The WhatsApp Case Study: examples	191

Chapter 1

Introduction

Our personality traits are expressed in our social relationships and daily activities and, naturally, are registered in the digital world, mainly in social media as we interact with each other on the various existing platforms [145, 68, 128]. Indeed, the pervasive use of such platforms has changed the way people interact with their peers, generating massive quantities of online data records. Such expressive volume of data capturing different facets of human behavior has motivated many research efforts to analyze and model psychological aspects behind the observed behavioral patterns [183, 146, 126, 117, 188]. For example, the rising of the Big Five personality test [58] as a powerful tool for measuring and characterizing the human personality traits has opened new opportunities for computing researchers around the world [128, 158]. By exploring models of human personality traits, such as those used as components of the Big Five test [183, 68], it has been possible to optimize various services and technologies for particular types of human personalities, ultimately increasing user satisfaction [126, 105]. Examples of such efforts are the use of personality models to improve search [102], recommender systems [186, 84], matchmaking services [68], educational technologies, [172] as well as mobile devices [156] and services based on the Internet of Things (IoT) [5].

Several studies of human personality have already shown that some types of personality traits are strongly associated with human curiosity [116, 128, 68, 183]. Indeed, curiosity has been recognized as a critical factor that influences human behavior in positive and negative ways at all stages of life: from a driving force in child development to the main driver of all basic scientific and philosophical inquiries. Curiosity has also been linked to many behavior disorders and non-sanctioned behaviors, such as use of drugs and criminal activity [54, 110, 83]. Similarly, curiosity has also been seen as an intrinsically motivated desire or appetite for information, or a passion for learning [83]. As such, the modeling and analysis of human curiosity as a driving force behind online user behavior deserves further investigation. This dissertation pursues this avenue of research, tackling the challenges of modeling and analyzing user curiosity in information services.

1.1 Motivation

As argued in prior work, an individual’s curiosity is essentially driven by external stimuli [16, 83] which are captured by various collative variables. These variables refer to different factors that govern how one’s curiosity is stimulated [16]. Novelty, complexity, uncertainty and conflict are examples of collative variables capturing aspects related to curiosity stimulation that have been reported in the literature [83]. According to Berlyne, novelty refers to how new the stimulus experienced is with respect to one’s past experience, whereas complexity refers to the degree of diversity in a stimulus [16]. Uncertainty, in turn, arises when an individual has difficulty in deciding how to respond to a stimulus, whereas conflict occurs when the same stimulus triggers two or more incompatible responses in an individual. All these factors ultimately affect curiosity arousal. As a matter of fact, one can argue that a stimulus is as a combination of multiple collative variables and its intensity is related to the arousal of curiosity in an individual.

In that direction, Wundt’s curve has been proposed to model, in general terms, the curiosity of an individual as a function of the intensity of the stimulus the individual is subjected to [177, 16]. This is a bell-shaped curve illustrating that as the stimulus intensity increases from low to moderate, its effect on curiosity is pleasant and rewarding (i.e., curiosity increases). However, as it increases to higher levels, its effects becomes unpleasant and even painful (i.e., curiosity decreases). In other words, too little stimulus leads to an individual state of boredom, whereas too much stimulus leads to one’s anxiety. Otherwise, there is an intermediate level of stimulus that leads to maximum curiosity. It is natural to model each individual with her own Wundt’s curve as the intensity thresholds defining the three zones (boredom, curiosity and anxiety) vary from person to person.

Curiosity models based on Wundt’s curve have already been applied in the design of artificial creativity systems, to control the behavior of non-player characters in digital games and in reconfigurable robots [47, 101, 45, 75]. In the particular domain of online information services, there have been recent efforts to formally introduce curiosity models into the design and evaluation of recommendation systems [186, 29, 179, 167, 139].

As an example, Zhao and Lee developed a recommendation system framework which considers that each recommended item denotes a stimulus to the user. Such stimulus, in turn, was related to curiosity scores using Wundt’s curve. The authors showed that the integration of such curiosity model into different recommendation mechanisms provides personalized recommendations with improved accuracy [186]. Yet, the authors have explored a single collative variable, namely novelty, as stimulus to curiosity and, as such, does not cover all components of human curiosity stimulation. Similarly, a few more recent efforts have also integrated curiosity models into service design, but exploring a single collative variable, often novelty [174, 111, 143, 1].

One notable exception is the work by Xu *et al.*, which explored two different collative variables, namely novelty and conflict [179]. However, the use of other collative variables, jointly or in isolation, remains to be investigated. Indeed, Wu *et al.* argued that computation models exploring psychological theories often consider just one (or a few) collative variable, disregarding yet how multiple collative variables affect the level of stimulation, individually or collectively. Nevertheless, the quantitative analysis of collative variables can help form a deeper understanding of the working mechanism of computational curiosity [171].

Most previous studies on stimulating curiosity in online information services have not only focused on a single collative variable, but have also neglected an important aspect that can influence curiosity, namely *social influence*. Social influence has been extensively studied as a key component in various behavioral phenomena, from information dissemination to opinion adoption (e.g., [12, 66]). Earlier studies of social influence can be broadly divided into two main categories: those that aimed to quantify social influence, often with the intention of identifying the most influential users in the system [152, 154, 57, 31, 2, 187], and those who proposed models of social influence diffusion [17, 79, 85, 55, 135, 103, 26, 49, 25, 50].

Our present effort is closer to the former group but has a key difference in purpose: we here aim at proposing metrics to quantify social influence as a component of an individual’s curiosity stimulation driving her towards behaving in certain ways, and further exploring such metrics to improve information services. To our knowledge, prior efforts to model social influence as a component of one’s curiosity stimulation are still quite scarce and recent [174, 170, 179]. Similarly, some few studies have explored principles derived from social influence in the design of recommendation systems [141], but without explicitly quantifying them as part of a curiosity model.

Despite such lack of prior efforts to explicitly model social influence as a component of curiosity stimulation, the general concept of *social curiosity* has already been widely discussed. Social curiosity, a facet of curiosity, has been defined as the general interest in gaining new social information motivating exploratory behaviors [52, 120, 80]. Specifically, social curiosity entails two different aspects: a general interest in obtaining new information about how others think, behave or act as well as an interest in interpersonal information that is obtained through exploratory behavior. For instance, social information is acquired through interpersonal conversations, i.e., through how one interprets daily events, hearing about each other’s lives and observing each other’s behaviors and expressions. As such, social influence can be seen as yet another collative variable that stimulates one’s curiosity and ultimately drives one’s behavior. Indeed Loewenstein [84, 83] asserts that an important reference point for individuals is the attainments of others, implying that an individual who adopts other people’s information sets as their own informational reference points become curious to acquire all the knowledge they possess.

In sum, efforts to operationalize social influence as a component of curiosity stimulation (i.e., as another collative variable) is a research avenue that remains mostly unexplored, offering the potential to push the state-of-the-art in curiosity modeling even further. In particular, one may expect social influence to play a particularly strong role in curiosity stimulation on social media platforms, where user interactions (e.g., posts, comments, likes) drive the dynamics of human behavior in general.

1.2 Problem Statement

In this dissertation, we define curiosity as the interest, attention, and time users devote to on their interactions (i.e., user sessions) in online information systems, such as information consumption (e.g., listening to songs, browsing the Internet, reading news articles) and information production (e.g., posting a message or picture, sharing video or website link), which are the two spectrums of the online information dissemination process. Thus, this dissertation aims at to investigate the following hypothesis:

Hypothesis *Models capturing multiple facets of human curiosity, such as social curiosity, can be used to uncover relevant user behavior profiles for the sake of enabling a more fundamental understanding of the online information dissemination process as well as designing more effective personalized information services.*

Towards addressing this hypothesis we aim at modeling and analyzing user curiosity as a driving force behind online behavior, notably information consumption and production/sharing. Building upon existing literature, we intend to advance the state-of-the-art by considering multiple collative variables as components of curiosity stimulation. In particular, we here argue for the importance of capturing social influence as a component of curiosity stimulation, notably in social media platforms. Moreover, we also investigate the use of Wundt’s curve as a model of user curiosity stimulation driving online information dissemination.

1.3 Objectives

We envision our investigation of the hypothesis posed in the previous section by tackling four complementary questions, as described next:

- **Research Question 1 (RQ1):** *Are there distinct user behavior profiles in terms of curiosity stimuli, as captured by multiple collative variables?* Towards answering this question, we aim at designing metrics capturing different collative variables associated with curiosity stimulation, notably those identified by Berlyne [16], that is, novelty, uncertainty, complexity and conflict. We also intend to use these metrics to characterize and uncover user behavior profiles that can help understand the online information dissemination process.
- **Research Question 2 (RQ2):** *How can we capture social influence as a component of human curiosity stimulation driving online information dissemination?* As mentioned, prior efforts to capture social influence as part of curiosity models have been very timid [174, 170, 179]. Yet, we speculate that such aspect may play an important role in the composition of the user curiosity driving the information dissemination process, at least for some users, particularly in the social media domain. Thus, towards tackling this question, we intend to cast social influence as another collative variable of curiosity stimulation, and design novel quantitative metrics to capture such variable. Moreover, we intend to show that this variable, as represented by the proposed metrics, captures aspects of curiosity that are not covered by other traditionally studied collative variables (e.g., novelty, uncertainty, etc). We then intend to use the proposed metrics to characterize user behavior and reveal relevant patterns and profiles driving information dissemination.
- **Research Question 3 (RQ3):** *To which extent user curiosity driving online information dissemination can be accurately modeled by a Wundt's curve?* This question is motivated by observations that human behavior (curiosity in particular) can be very dynamic and heterogeneous [83]. Such dynamics and heterogeneity have already been widely reported with respect to online behavior [81, 28, 166, 169]. As such, it is unclear whether Wundt's curve, which has been analyzed mostly from a theoretical perspective, can indeed be a reasonably accurate representation of human curiosity stimulation for real users of online information services. Indeed, to our knowledge, no prior work has investigated the adequacy of such model for capturing human curiosity stimulation driving information dissemination. This question is addressed by exploring alternative strategies for operationalizing the general ideas behind Wundt curve (i.e., two thresholds that delineate the areas of maximum cu-




Figure 1.1: Pictorial representation of research questions of dissertation.

riosity stimulation from boredom and anxiety) and evaluating the fit of the resulting models to different user profiles.

We intend to address the aforementioned research questions by exploring different case studies in the information dissemination domain. We illustrate how these four research questions complement each other, towards exploring our posed hypothesis, in Figure 1.1.

1.4 Contributions

The contributions achieved so far are concentrated in the first three research questions, i.e., RQ1, RQ2 and RQ3. Towards answering these questions, we explored two case studies related to information dissemination. In our first case study, focused on the LastFM platform [133], we investigate the role of human curiosity in the *consumption* of online information, specifically online music. Our second case study, in turn, tackles the other end of the information dissemination process, namely information *production*. To that end, the focus is shifted to a platform where social interactions are a fundamental driver for user behavior, namely WhatsApp groups.

Specifically, our contributions so far are:

- Proposal and analysis of seven metrics capturing different collative variables associated with curiosity stimulation in the context of information consumption (RQ1). As mentioned, most prior work has focused on a single collative variable, notably novelty (and its three main components, recency, frequency and dissimilarity). We here take a step further by analyzing different variables (novelty, uncertainty, complexity and conflict) and investigating the relationships among them in terms their redundancy or complementarity in capturing aspects of curiosity stimulation. We proposed metrics to the particular scenario of online music consumption (first case study) and use them to uncover different profiles of user curiosity stimulation in LastFM.

- Proposal and analysis of novel metrics capturing social influence as a collative variable associated with human curiosity stimulation (RQ2). The metrics were designed to capture curiosity stimulation as a driver behind information production. Moreover, we focus on an communication environment, i.e., WhatsApp, where user behavior is expected to be mostly driven by their social interactions while sharing and exchanging messages (second case study). Specifically, we propose five metrics of social influence at an individual level as well as one metric capturing social influence at the (WhatsApp) group level. We then use the metrics to show that social influence is indeed complementary to other (priorly studied) collative variables of novelty, conflict, complexity and uncertainty. In other words, our metrics are able to capture aspects of curiosity stimulation not covered by the traditionally studied variables, thus reflecting a novel and important component of the curiosity stimulation process. Moreover, we also use the metrics to offer an extensive characterization of user-level and group-level of curiosity stimulation.
- We model and analyze curiosity stimulation, as a driver behind online information dissemination, by focusing on the two ends of the process, that is, information production (and sharing) and information consumption, in two complementary case studies. In our first case study, we analyze curiosity stimulation as a driver for online information consumption, specifically music content. To that end, we use a large dataset containing records of user listenings to different songs on LastFM, one of the largest online music platforms and currently is a personal metadata repository of user's historical listenings [114, 133]. In our second case study, in turn, we focus on the information production and sharing end. Specifically, we delve into the role of social influence as a force that stimulate user's curiosity to share content. To that end, we focus on a currently very popular communication platform, notably WhatsApp, and look into the curiosity as a driver of user sharing content in different groups.
- Adaptations of the general Wundt's curve as models of curiosity stimulation to capture the diversity and complexity of mixed patterns observed in real users (RQ3). Based on extensive analysis of real datasets capturing user listenings to different pieces of music (our first case study), we find that the curiosity stimulation curve for a large fraction of users is indeed multimodal, suggesting much more complex curiosity stimulation patterns than those expressed by the simple Wundt's curve. Nevertheless, we also find that, for most of such cases, a combination of two or three Wundt's curve is able to capture reasonably well such mixed patterns as well as the dynamics of user curiosity over time. Furthermore, we use the metrics of social curiosity together with the traditional collative variables on a different platform to investigate the aspects of arousal curiosity in the production and sharing of

information in group communication (second case study). In this case, we find that the stimulation of social curiosity plays an important role in estimating the curiosity curve because it increases the number of users that can be accurately modeled by a Wundt's curve. We also find that there are cases that require more than one Wundt's curve, which proves that users have complex behavior patterns even in a content sharing environment.

1.5 Outline

The rest of this dissertation is organized as follows: Chapter 2 discusses background and prior work in areas closely related to the topic of this dissertation. Chapter 3 presents our problem statement along with key assumptions and associated notations. Chapters 4 and 5 present our exploration of the three research questions, namely RQ1, RQ2 and RQ3, in two different case studies. Finally, Chapter 6 provides our final considerations, summarizes the contributions and concludes this dissertation, as well as presenting future research directions.

Chapter 2

Background and Related Work

In this chapter, we present a summary of background knowledge and related work that are essential to the understanding of this dissertation. Figure 2.1 shows a mind mapping of the the main domains to which the research developed in this dissertation relates. In light of this mapping, this Chapter is organized as follows:

- Section 2.1 provides the main concepts and theories of Psychology of curiosity, notably different dimensions of curiosity, the information gap theory and its relationship with optimal level of curiosity stimulation and social curiosity;
- Section 2.2 presents computational curiosity models, notably models used in on-line information systems, and discusses prior work on different applications of such models, including recommendations systems, which are more closely related to the scope of this dissertation;
- Section 2.3 reviews prior work on social influence and its important role in the information dissemination process;




Figure 2.1: Mind mapping of the domains which compose this dissertation.

2.1 Psychology of Curiosity

Curiosity is a powerful force that drives our daily occupations [96]. In particular, human beings dedicate a lot of time in activities related to information seeking and consumption (e.g. browsing the Internet) [65]. All these activities are driven by our need for information, that is, by *curiosity*. Curiosity has been recognized as the desire for information which facilitates learning, promote new discoveries and enriches life of knowledge [116]. In contrast, Hsee and Ruan highlight that the desire to resolve curiosity can lead humans to seek information despite predictably undesirable consequences [54]. This means that they do so to satisfy their curiosity without paying attention to potential negative consequences. Yet, curiosity has been seen as an intense, intrinsically motivated appetite for information, that is, an experience analogous to physiological appetites which may produce painful feelings of deprivation if not satisfied [83].

Curiosity is not singular, it has many forms and can direct itself in a number of directions [140]. Berlyne has defined two dimensions of curiosity: *perceptual* curiosity and *epistemic* curiosity [16]. The former relates to a driving force evoked by a novel stimuli and is reduced by continuous contact with these stimuli, such as, for example, the exploratory behavior in animals [16, 83]. *Epistemic* curiosity, in turn, concerns a desire for knowledge and is generally associated to human beings. This is the type of curiosity that is arisen by complex ideas or conceptual ambiguities (e.g. scientific theories, intellectual enigmas), which motivate us to ask questions or test hypotheses in order to acquire knowledge [30]. Another commonly studied characterization of curiosity has been that distinguishing between *state*, which is curiosity in a particular situation, and *trait*, which is individual differences in capacity to experience curiosity [83, 21]. Our focus in this dissertation is on the *state* of *epistemic curiosity* of human beings and its role in online information dissemination.

According to the Information Gap Theory [83], one's curiosity arises due to a discrepancy between *what one knows*, fairly objective, and *what one wishes to know*, highly subjective. This implies that an individual should have awareness of information gap for experiencing curiosity [10, 42]. Thus, curiosity should be positively related to one's knowledge in a particular domain [43]. According as [16], one's curiosity is induced by external stimuli, more specifically "stimulus conflict" or "incongruity". This involves properties such as complexity, novelty, and surprise. So, the arousal of curiosity on an individual depends on the appropriate level of stimulation that can be induced on her by a given stimulus.

Wundt introduced the concept of *optimal level of stimulation* and postulated a bell-shaped relationship between the stimulation level and the hedonic response [177]. This hedonic response is directly associated with curiosity as it represents the level of




Figure 2.2: Wundt's curve and its zones of curiosity stimulation.

pleasure one feels while satisfying her curiosity triggered by the input stimulus. The optimal level of stimulation is also called arousal potential [16]. An example of Wundt's curve is shown in Figure 2.2, where the axis x denotes the level of stimulus received by an individual and the axis y denotes the individual's hedonic response to such stimulus. Lower hedonic values (unpleasantness) are achieved if the user receives too much (i.e. anxiety zone) or too little (i.e. boredom zone) stimulation, whereas maximum hedonic value (thus maximum curiosity) is triggered by moderate stimulus (i.e. *curiosity zone* or *arousal tonus*). The two ends of spectrum in Wundt's curve reflect two rules in stimulus selection: *avoidance of boredom*, which is related to the boredom zone, and *avoidance of anxiety* regarding the anxiety zone [171]. The intermediate region determines the *optimal level of stimulation*. Each individual has her own Wundt's curve as the thresholds defining the boredom, curiosity and anxiety zones vary across different people.

According to [16], different *collative variables* govern the curiosity stimulation process, notably:

- *Novelty*, occurs when a stimulus is new to an organism and is related to the time at which the organism experienced the stimulus. Intuitively, it is inversely related to (i) *frequency* in the sense of how often the stimulus were previously experienced, (ii) *recency* in the sense of how recently the stimulus were experienced, and (iii) *similarity* in the sense of how similar the stimulus are to previously experienced stimulus. For example, a stimulus can be completely new to an organism, that is, it has never been experienced before, or it can have been experienced in the past in a short period of time. There are three meanings of novelty: complete novelty, when something is completely new to the organism; short-term novelty, when the new sensation was only present to the organism for a short time (e.g., for a few minutes); and long-term novelty, when something is new to the organism over a period of several days. There are also additional variables to novelty: (a) *change* refers to a movement when the stimulus acts on the individual's receptors (i.e., a change in the stimulus) and may cause some priority in exploratory directions;

- (b) *surprising* arises when there is a stimulus that contradicts the expectation; (c) *incongruity* indicates an expectation that is not met by the same stimulus. As an example of novelty, let us consider the first appearance of a smartphone in the market or the discovery of a new biological species [82];
- *Uncertainty*, which is related to the difficulty of organism in deciding how to respond to a stimulus. To understand this concept, we need some context. In psychology, there are two information spaces, the stimulus space and the response space, which are of interest and can be influenced by the type of stimulus that occurs and what responses may occur instead and how likely they were. So, to define the stimulus uncertainty it is needed to describe the proper information spaces. This means that (i) a list of stimuli that can occur should be drawn up, (ii) these should be divided into classes and (iii) a probability should be defined for each class. In this sense, the probability of each class indicates the competing strength of possible responses and the degree of uncertainty is denoted by the *entropy*, i.e., the quantity of information of stimulus space (see Section 4.2.2 for more details). Therefore uncertainty describe situations in which any number of alternative results is possible. For example, the weather forecasts is often related to uncertainty, notwithstanding, with the aid of modern computer models of advanced technology, the meteorologists yet sometimes wrong their predictions [82];
 - *Conflict*, occurs when a stimulus causes two or more incompatible responses in an organism (e.g., moving forwards and backwards at the same time or raising and lowering a leg at the same time). A response may be incompatible in several ways: (a) they may be inherently contradictory; (b) other responses may initially be performed together but become incompatible through learning; (c) limiting an organism's ability to multitask. Complementarily, conflict is positively related to (i) the approximate equality of the strengths of the competing responses (i.e., the probability of each stimulus class), (ii) the absolute strength of the competing response, (iii) the number of competing responses, and (iv) the degree of incompatibility between the competing responses. In summary, conflicts characterize situations in which new information is incompatible with existing knowledge;
 - *Complexity*, refers to the degree of variety or diversity in a stimulus pattern. It is assumed that the degree of complexity is positively related to: (i) the number of distinguishable elements in a stimulus; (ii) the degree of dissimilarity between elements; (iii) the extent to which many elements are perceived and responded to as a unit. For example, a figure consisting of many geometric shapes is generally considered more complex than a figure consisting of only one simple geometric shape.

Those are called collative variables because their evaluation involves an analysis of

similarities and differences between existing elements in a stimulus pattern. In general, these variables refer to different external factors that govern how curiosity is stimulated. Thus, a stimulus is indeed a combination of multiple collative variables.

In his seminal work, [16] argues that properties of such collative variables can be discussed in an information-theory related language, and proposes a methodology to quantify such variables based on information theoretical metrics. Following a similar argument, [144] states that, with some mathematical manipulation, all collative variables associated with curiosity stimulation can be expressed by formulations using metrics from information theory. As we discuss in the next section, these arguments inspired some recent studies [186, 179, 6, 174, 170, 143, 1] in Computer Science to propose metrics capturing different collative variables, notably those originally investigated by Berlyne (i.e., novelty, conflict, uncertainty and complexity).

We here argue that, in addition to the four aforementioned collative variables, *social influence* should also be considered as part of the stimulus to one's curiosity, especially in online social networks and social media applications, where, to a large extent, one's behavior is driven by connections, common interests with others and social observation.

Recently, [63] proposed a five-dimensional model of curiosity, which can be quoted: (1) joyous exploration, (2) sensitivity to deprivation, (3) stress tolerance, (4) thrill seeking and, most importantly for this thesis, (5) social curiosity. These dimensions are interrelated, but they are also independent of each other and can be distinguished by links to personality, emotions and well-being. According to the authors, social curiosity denotes the individual skills to tackle the interpersonal world, that is, the arising of interest in obtaining new information about how others think, behave or act as well as an interest in interpersonal information which is obtained through exploratory behavior [63]. Our argument is aligned with the concept of *social curiosity* which is an essential component of the five-dimensional curiosity model.

Others have also offered definitions and perspectives on the role of social curiosity on human behavior. For instance, [52] define social curiosity as an interest in gathering new social information motivating exploratory behaviors, whereas [80] argue that social curiosity includes individual's experiences and states in the social world (e.g., for social comparisons, for forming friendships and for attacking adversaries). [120], in turn, declares that social curiosity is an essential element for building (and using) social networks and relationships, which, in turn, is a central human task. More recently, [61] distinguish between *overt* and *covert* social curiosity. The former refers to being curious about how others think, behave or act, and is related to healthier outcomes (e.g., interpersonal competencies). The latter refers to observing others in a indirect and secretive way to obtain new information and regards to unhealthy outcomes (e.g., gossiping and social anxiety). Despite all such prior studies, to our knowledge, efforts to *quantify* social curiosity, notably in the online world, are quite scarce.

In general, curiosity plays an increasingly large role in the economics of information, especially because modern technology is geared towards attracting, maintaining and redirecting user attention [168]. Therefore, it is arguably one of the dominant forces driving many sources of economic growth and educational attainment. Indeed, thanks to the unprecedented availability of large amounts of data, there has been a recent surge of studies on psychological aspects of online behavior, notably curiosity and personality traits [145, 78, 71, 183, 146, 126, 117, 188].

By characterizing one's curiosity and personality traits, it has been possible to tune and enhance services and technologies (e.g., search and recommendation services, educational and mobile technologies) to better meet one's individual personality, ultimately improving user satisfaction [186, 5, 84, 172]. We further elaborate on such efforts in the next section.

2.2 Computational Models of Curiosity

Most existing computational models of curiosity are grounded in a general appraisal process that can be abstracted into a uniform two-step model as follows [171]:

- **Step 1:** Evaluation of the *stimulation level* based on collative variables;
- **Step 2:** Evaluation of the *curiosity level* based on the principle of intermediate arousal potential, i.e., the optimal level of stimulation.

According to the aforementioned general appraisal process, the input data is used to estimate the stimulation level, based on (one or a combination of) collative variables, as defined by the computational model. Next, the inferred stimulation level is mapped onto a curiosity level following, for instance, the *optimal level of stimulation* and Wundt's curve discussed in Section 2.1. Finally, the curiosity level (and at times, even the stimulus level directly) are then used as component of a computational application.

In general, curiosity models have been broadly utilized in several computational areas, from artificial intelligence and education to online information services. Figure 2.3 provides an overview of various applications of computational curiosity models in several areas of Computer Science. As shown in the figure, the most important areas [171] are artificial intelligence, artificial creativity, games, education and online information systems.

The most frequent use of curiosity models is, by large, in *artificial intelligence* [150], including machine learning, natural language processing, intelligent agents and robotics.




Figure 2.3: Applications of computational curiosity models.

In machine learning, curiosity models have been used in the design models to tackle self-supervised learning problems such as visual categorization task [159]. In natural language processing, for example, those models have been employed in the automatic description of images using curiosity-driven reinforcement learning for visual paragraph generation [87]. Other uses of curiosity model in artificial intelligence include the design of agents that can explore environments [90, 45, 75], the autonomous social robots that interact with children in educational settings [44], as well as interactive and adaptive learning system for shopkeeper robot [36], among others [24, 13].

In the *artificial creativity* area, it is worth mentioning the seminal work of [129] describing an approach to developing computational models of creativity using agents that incorporate a hedonic function based on Wundt Curve. Also, [93] utilize a creativity behavior model with a combination of curiosity and learning in an agent-based system to focus attention of learning agents. [47], in turn, proposed a framework where a creative systems can intentionally exhibit transformational creativity, that is, it captures the notions of unexpectedness, surprise and specific curiosity.

In *games*, [101] proposed the use of reinforcement learning together with psycholog-

ically inspired motivation models for creating curious non-player characters in multi-user games. In another approach, [130] used the identification of curiosity as player’s motivation to improve game design.

In *education*, curiosity models based on the collative variables of novelty, surprise and uncertainty have been explored in the design of pedagogical agents (i.e., tutors or companions) to improve human learning in virtual environments [172]. Similarly, [173] explored curiosity models to simulate learning companions that use knowledge of users and environments to identify interesting objects. Also, [175] proposed a curious peer learner in a virtual learning environment to simulate the student’s behavior. Similarly, [74] proposed a learning by teaching platform in which students work individually or in groups to teach a conversational agent how to classify objects. In a different perspective, [119] designed an approach to aid a curiosity-inspired learning in every day life.

Finally, in the context of *online information systems*, curiosity models have been employed in online advertising [162], online social networks [126], crowdsourcing systems [73], recommendations systems [186], search systems [102] and online purchase [9]. As these applications are closer to our present focus on online information dissemination, we further elaborate on prior efforts in that domain in the following sections.

2.2.1 Curiosity Models in Online Information Systems

In this section, we focus on applications of curiosity models in *online information systems* (blue branch at the bottom of Figure 2.3). As shown in the figure, those models have been employed in various subareas of interest.

For instance, in the particular context of online advertising, a seminal work is that of [100], in which the authors studied the power of curiosity arousal in online advertising by tracking click-stream variables (e.g., records of screens visited and exposure time) as well as behavioral intention data. More recently, [162] proposed the use of curiosity in a clickbait detector by assessing the curiosity stimulus on readers exposed to several types of news headlines.

Vega-Oliveros *et al.* [161], in turn, incorporated the notion of social curiosity into epidemiological models of information diffusion in online social networks. In particular, the authors considered that each user’s curiosity is unique and, unlike most prior studies, dynamically evolves over time. [126], in turn, investigated the prediction of a user’s curiosity based on features gathered from Facebook accounts. In a different context, [9] analyzed the role of curiosity in user’s purchase intention or buying in e-commerce websites. To that end, they created a survey to evaluate the involvement, engagement

and satisfaction of user and how much its behavior is influenced by curiosity.

Millan-Cifuentes *et al.* [102] studied online search by curiosity, that is, search driven by an hedonistic need, rather than by a specific information need. To that end, the authors designed a social media search application in which the user's search is triggered by her spatial-temporal context (and not by the traditional query-response based information retrieval). The authors simulate a casual-leisure search scenario in which the user's context serves as stimulus to arise her curiosity and trigger the search.

There has also been efforts to study curiosity in crowdsourcing systems. For instance, [73] studied the potential of curiosity as a novel kind of intrinsic motivation to encourage crowd workers. The authors grounded their study on *the information gap theory* and applied curiosity-inducing designs in crowdsourcing task interfaces to improve engagement and performance.

The aforementioned studies offer a sample of recent efforts to apply curiosity models in the study of online information systems. In this dissertation, we focus on employing such models to study information consumption and production in social media applications. Additionally we also emphasize that their use to improve user experience in recommendation systems, yet another area of application that has been explored in recent literature. Given such close application connected to our present effort, we delve deeper into those efforts in the next section.

2.2.2 Principles of Curiosity Models from Recommendation Systems

As we saw at the beginning of this section, a lot of research is being done on curiosity models in the field of artificial intelligence and machine learning, which is why these models are better established than other areas [150]. Since the modern techniques and algorithms from this field have a wide range of applications in online information systems, curiosity models of recommendations systems are by far the most widely used and most developed [171]. Therefore, these curiosity models, which make use of the two steps of the general appraisal process, form the most important foundations for achieving advantages in other areas.

In this context, prior efforts to introduce curiosity models in recommendation systems can be generally grouped into two large sets depending on how user curiosity is inferred. One set of studies incorporates curiosity into the recommendation models by manual intervention through the gathering of user questionnaires to learn their curiosity levels. Another group of studies, closest to this dissertation, infer user curiosity from

behavioral data (e.g., user traces) by first employing specific metrics to quantify collative variables based on the input data and then incorporating such metrics (and inferred curiosity) directly into the recommendation process. Below, we review previous studies in each category.

2.2.2.1 Inferring User Curiosity from Questionnaire Data

Santos proposed a hybrid recommendation system that considers the curiosity level of each user as a factor in site recommendation [125]. User curiosity is inferred based on questionnaire data gathered by a survey test. Later, the same author and colleagues proposed CURUMIM, an online system that generates serendipitous, personalized and novel recommendations for tourist places [127]. The curiosity profile model is intrinsically based on informational data from online social network which is highly correlated with items of a survey test.

Maccatrozzo *et al.* [88, 89], in turn, proposed a model, called SIRUP, to operationalize user curiosity in content-based recommendation. Inspired by Silvia's [144] and Berlyne's [16] theories on curiosity stimulation, discussed in the previous section, the authors focused on one particular collative variable, namely *novelty*. The authors used a metric to quantify the novelty of items in a TV program, applied a questionnaire to assess how users coped with *novelty*, and used the survey data to infer user curiosity profiles. They then used such profiles to create personalized recommendations.

Chen *et al.* [29] conducted a large-scale user survey in Mobile Taobao, a popular app of e-commerce in China, to measure user perception regarding serendipity, novelty, diversity, user satisfaction and intention of purchase. The authors also concluded that serendipity-oriented recommendation algorithms produce better results, as perceived by users and reported by them in the survey, compared to relevance- and novelty-oriented strategies. The same authors also analyzed the correlations between different item features and user characteristics (e.g., personality traits) the user perceptions of serendipity in recommendations, offering insights into how to improve serendipity-oriented recommender systems [167].

In a different direction, [104], in turn, proposed Pique, a personalized method to recommend sequences of scientific papers to students customized to their background and interest aiming at encouraging curiosity.

In common, all aforementioned studies relied on user surveys and questionnaires to infer user curiosity and its effect on recommendation. In this dissertation, we aim

at proposing metrics to quantify different collative variables associated with curiosity stimulation from behavioral data (e.g., traces of usage of particular applications). As such, our work is more closely related to those discussed in the next section.

2.2.2.2 Inferring User Curiosity from Behavioral Data

One of the earliest efforts to infer user curiosity from behavioral data to support the design of a recommendation system is that by [186]. The authors proposed the *Curiosity-based Recommendation System* (CBRS) framework to generate personalized recommendations by explicitly modeling the curiosity stimulation curve of each user, i.e., the user’s Wundt’s curve, from the user’s access history. To do so, the authors introduced metrics to quantify the curiosity stimulated on a user by a given item based on three aspects associated with *novelty* (one of the collative variables introduced in [16]), namely *frequency*, *recency* and *dissimilarity*. These metrics were used as input to a probability distribution function mapping stimulus to curiosity score, approximated by a Beta distribution. The curiosity score, in turn, was then used as component of the recommendation function. In a similar strategy, [143] also explored metrics associated with the same three aspects of *novelty*. But unlike [186], the authors used a Normal distribution to model the user’s curiosity curve. Based on these curves, they proposed the *Curiosity Inspired Personalized Recommendation Re-ranking Framework* (CIR) aiming tackle overcome the trade-off between recommendation accuracy and novelty.

Abbas and Niu [1] also proposed a recommender system framework grounded on the use of novelty as a source of curiosity stimulation. Yet, unlike [186, 143], the authors explored metrics related to *frequency* and *surprise* to quantify stimulus intensity as a form of *novelty*. *Surprise* in turn occurs when a stimulus triggers a certain expectation and then another stimulus occurs that contradicts this expectation. So, that is, *surprise* together with *change* and *incongruence*, are supplementary variables for *novelty* [16]. Also, they approximated Wundt’s curve of each user by a Beta distribution, parameterized according to the input data.

More recently, [179] proposed the *Curiosity-drive Recommendation Framework* (CdRF) which also incorporates a curiosity stimulation mechanism into the recommendation process. To that end, the authors explored metrics associated with two collative variables, namely *novelty* and *conflict*. In particular, the authors proposed metrics to quantify *social conflict*, defined as the closeness between positive and negative responses (i.e., ratings) on the same item received from the user’s close social peers. They also modeled the Wundt’s curve of each user as a predictive curiosity function which maps a

stimulus intensity (estimated by a combination of the proposed metrics) into a curiosity score, which, in turn, is used as a component of the recommendation function.

What all of the above studies have in common is that they proposed metrics to quantify curiosity stimuli, used them as input for explicit models of the user’s curiosity stimulus curve (inspired by Wundt’s curve) and in turn used these models as part of the recommendation process. In other words, they implement the two steps of the general appraisal process, namely the evaluation of the *stimulus level* and the evaluation of the *curiosity level*. In contrast, some other authors have opted to introduce metrics related to the stimulation of curiosity directly into recommender systems, thus avoiding the mapping from stimulus level to curiosity level, i.e. avoiding the step of evaluation of *curiosity level* in the general appraisal process. In that direction, some studies focused on *surprise* (again a supplementary variable of *novelty*) as the main driver of curiosity stimulation [6, 111]. For example, in [6], the authors explored alternative approaches to introduce *surprise* into a recommender system: one based on statistical co-occurrence likelihood and another based on user feedback. In contrast, the authors of [111] explored an existing computational model of *surprise*, aiming at improving satisfaction and inspiring curiosity of users of a health news recommender. In addition, [39] have developed a recommendation system that uses a single Beta distribution to model Wundt’s curve and reorder the recommendation items together with current and state-of-the-art Deep Learning recommendation algorithms by using a metric of *surprise* based on the pointwise mutual information between item categories. Compared to previous works that use the general appraisal process based on Wundt’s curve, such as [186, 143, 1, 179], the authors rank the items by calculating the distance of the stimulus items in relation to the curiosity zone, and they do not use the curiosity score.

Surprise was also the main concept explored by [174], though in a somewhat different context. Specifically, the authors proposed a recommendation model grounded on *social curiosity* based on the idea that unexpected ratings of friends with respect to a given item triggers the user’s curiosity towards the item. The proposed model combines user preferences (i.e., predicted ratings) and curiosity (as captured by the surprise with respect to items), taking into account the historical ratings of the user and of his friends. Predicted ratings and curiosity scores are then combined into an interest score used to rank the recommended items. In a follow-up study, the same authors turned their focus to another dimension of curiosity stimulation, that is, *uncertainty* [170]. They modeled the feeling of *uncertainty* elicited on the user when his/her friends give different ratings to the same item. They then proposed a personalized recommendation strategy based on the aggregation of the rankings of items based on user preferences and on curiosity, here estimated with respect to *uncertainty*.

Table 2.1 offers a summary of the aforementioned efforts to introduce curiosity on recommendation systems by exploring behavioral data. For each study, the table lists the

key components of the two steps of the general appraisal process of curiosity modeling: (i) *Step 1* is the evaluation of the stimulation level based on collative variables; and (ii) *Step 2* is the estimate of the curiosity level as a function of the stimulation level. For the latter, the last column, if applicable, indicates the type of curve of the curiosity profile, which in most cases is a *unimodal curve*, generally resembling a bell-shaped curve. We emphasize that a *unimodal curve* is not able to capture different aspects of stimulus patterns and long-term changes, as the user's curiosity can be stimulated in different ways in different situations and can evolve over time. For this purpose, a *multimodal curve* is required, which we will discuss in the next chapters. In particular, note that all those studies considered either a single or at most two collative variables associated with curiosity stimulation. One contribution of this dissertation is to propose metrics to capture a larger set of collative variables and assess to which extent these variables complement each other as component of the curiosity stimulation process. We do so for two different case studies related to online information consumption and production, discussed in Chapters 4 and 5, respectively.

Table 2.1: Summary of prior efforts to introduce curiosity into recommendation systems based on *behavioral data*.

Reference	General Appraisal Process		Modeling of Curiosity Profile
	Collative Variable (Step 1)	Stimulus Selection Process (Step 2)	
[186]	novelty	Wundt's curve using <i>Beta distribution</i>	unimodal curve
[143]	novelty	Wundt's curve using <i>Normal distribution</i>	unimodal curve
[1]	novelty	Wundt's curve using <i>Beta distribution</i>	unimodal curve
[179]	novelty and social conflict	Wundt's curve using <i>histogram function</i> and <i>sigmoid functions</i>	unimodal curve
[6]	surprise	stimulus intensity	–
[111]	surprise	stimulus intensity	–
[174]	surprise	linear combination of user preference and stimulus intensity	–
[39]	surprise	Wundt's curve using <i>Beta distribution</i>	unimodal curve
[170]	uncertainty	linear combination of user preference and stimulus intensity	–

2.3 Social Influence

In the previous section, we discussed prior efforts to model curiosity in computational applications, notably recommender systems. As argued, these studies are grounded on traditionally studied collative variables associated with curiosity stimulation, most often novelty or surprise (which is a supplementary aspect of novelty). Yet, one of the hypotheses investigated in this dissertation is that *social influence* may also be a component of (social) curiosity stimulation, notably in social media applications where user actions are often influenced by the peers. Therefore prior studies of social influence constitute an important related area of exploration, briefly reviewed in this section.

To our knowledge, prior efforts to model social influence as (part of) a stimulus to user curiosity, as is one of our present goals, are quite rare in the literature. Yet, there is a rich body of prior analyses of social influence in social media platforms from different (though related) perspectives [151, 77, 51, 123, 124]. Those studies can be broadly grouped into two major collections: those that aimed at quantifying social influence, often aiming at identifying the most influential users on the system [152, 154, 57, 31], and those that proposed models of social influence diffusion (such as epidemic models [17], cascade models [79, 85], linear threshold model [55, 135] and complex contagious models [103, 26, 49, 25, 50]). One of the contributions of this dissertation is to propose metrics to quantify social influence as a component of curiosity stimulation, a topic that is closer to the former group. Thus, we review studies in that category in this section.

Some early attempts to quantify social influence relied only on network properties such as degree distributions, diameter, clustering coefficient, community aspects and small world effect, often considering a single network model capturing explicit connections among users (e.g., friendship links) [153, 33, 70, 51]. However, the limitations of considering only topological measures to estimate user influence have already been pointed out by prior studies [185, 163]. For instance, [185] argued that node degree (e.g, number of followers on Twitter) cannot be used alone to assess user influence as users with more connections (i.e., larger degrees) may not be the ones who more often forward content. [163] also argued that structural measures of influence can lead to misleading assessments of influence. For instance, higher popularity does not imply necessarily higher influence. Moreover, topological properties are inherently dynamic [38], which implies that their temporal evolution, which cannot be captured by a single overall network structure, must be considered.

Also, most prior studies assume that user influence propagates over explicit links that connect users on the platform under analysis (e.g., friendship links) [108, 109, 176]. However, in some platforms such as group communication applications (e.g., WhatsApp, Telegram), no such explicit links exist. Similarly, some prior analyses of user influence

relied on explicit models such as epidemic models [148], topic models [153], probabilistic graphic models [155] and statistical models [18]. Instead, our present effort to model social influence as a component of curiosity adopts a more general *model independent* approach, grounded on information theory metrics. Indeed the metrics we propose are measured in bits (is the unit of information which is 0 or 1, it is also called of *shannon* in information theory), which facilitates straightforward comparisons between systems [157]. Moreover, prior model-based studies of user influence often relied on linear models [46, 153]. However, online social networks are known to present *non-linear* relationships influencing information spread [33, 97]. In that sense, our information-theoretic metrics are more robust as they are capable of capturing linear and *non-linear* interactions [157, 136, 142].

Another body of studies relied on statistical or machine learning models to estimate user influence. For example, [46] proposed models of probabilistic influence among users and developed algorithms for learning the parameters of these models. [86], in turn, proposed a deep neural network based model to estimate social influence by exploiting the history of user actions propagated among friends. [76] presented measures to capture the social influence at both microscopic (based on explicit user actions such as comments and mentions) and macroscopic (based on number of followers) levels, while [184] used regression techniques to study the role of triadic patterns in user behavior prediction. [69], in turn, analyzed the probability of an individual to be socially influenced based on several machine learning approaches and a rich set of features including personal network exposure, structural diversity, locality, retweet time delay as well as size and path length of cascades.

In contrast to the aforementioned studies, [18] adopted a *causal* approach to the analysis of social influence from propagation data. Their goal – retrieving a minimal causal topology from the data – is somewhat complementary to ours. Instead, we here want to estimate how a user’s curiosity may be stimulated by others in a group communication platform. Other studies of social influence relied on topic modeling to assess the social influence between users in specific topics [155] and factor graphs, a probabilistic graphic model to learn the effects of social influence, action correlation between users and time-dependency of user actions [153].

Finally, the work by [163, 164] is arguably the closest to our present effort. They also proposed to estimate social influence based on information theoretical measures, but they used a different metric – information transfer¹, which is based on mutual information. They used information transfer to measure the *direct* social influence between users on Twitter and identify influential users based on their capacity to predict the behavior of

¹Information transfer is also called of transfer entropy, information transfer between two stochastic process characterizes the reduction of uncertainty in a process due to knowledge of the other process [136].

other users. Though with a similar goal, our work described in Chapter 5 differs from Steeg *et al.*'s by considering a completely different platform – WhatsApp, using somewhat different metrics, and considering both direct and indirect social influence as potential sources of curiosity stimulation.

More broadly, our work presented in Chapter 5 has some key differences in purpose from the aforementioned prior efforts, and, as such, has a complementary goal compared to them. Whereas most prior studies discussed above aimed at quantifying the influence of a user on others in general, often over a particular time window and focusing on the user who influences the others (origin of influence), we here aim at estimating how social influence from others may be driving a user towards sharing content (i.e., focus on the destination of influence). As such, social influence has to be computed at much finer granularity, i.e., each time a user shares a piece of content, since human curiosity is considered highly dynamic and contextual (see discussion in Section 3.2).

Aiming at distinguishing our work on social influence from prior studies, Table 2.2 presents a direct comparison of them along several dimensions. Some of these dimensions, such as whether the proposed solution is based on a *non-linear method* and whether it is *model independent*, refer to the level of generality of the proposed approach. For example, as mentioned, a non-linear model, as ours, should be more robust to capture non-linear effects impacting information diffusion in online social networks [97]. Similarly, by not exploring any specific model, as done by others [153, 148, 155, 18], we offer a more general approach, where the result is not a model parameter but rather a number quantifying some relationships that exist in the data [157], and free from particular assumptions and specificities that may constrain such models. Also, our study, unlike some others, does explicitly consider the *temporal dynamics* of social-influence driven curiosity, and by doing so, recognizes that this is a highly transient, dynamic and contextual human behavior trait [83, 61, 72].

Other dimensions listed in Table 2.2 relate to the particular domain of study – WhatsApp groups. As the first to propose metrics of curiosity stimulation to such environment, we had to address several challenges. For example, a *network free approach* was chosen due to the absence of explicit links connecting users. Instead, all group members may interact with each other at any time, unlike in other setups where there are explicit links connecting users through which social influence may propagate. This implies that, in the absence of such explicit links, a strategy to capture the heterogeneity and the dynamics of social influence in our metrics is required. Moreover, by developing a network free approach, we offer a more general solution that does not rely on existing links to estimate social influence and thus may be adapted to other setups.

Similarly, the target environment motivates the development of metrics to capture other effects of social influence, in addition to the often studied direct influence of one user on others. On one hand, we propose a metric to estimate *social curiosity at the*

group level, which can be used to characterize collective behavior in different groups. Since WhatsApp allows only for small groups (at most 256 members simultaneously), and these often target specific topics of discussion (as defined in the group name or description) [23, 122], characterizing the curiosity of such small user populations over time may offer useful insights into social behavior (much more than in more open and unconstrained spaces like Twitter and Facebook). Moreover, inspired by prior analyses of WhatsApp [22, 113], we also propose metrics to explicitly capture the *indirect effect* of strong influencers that may emerge in such constrained spaces.

In sum, compared to the existing literature, we here propose to quantify social influence from a novel perspective – curiosity stimulation – offering a more general approach that is inspired by arguments available in the Psychology literature [16, 144], and that addresses particular challenges of a currently very popular communication platform, namely WhatsApp groups.

Table 2.2: Overview of prior work and our present effort to estimate social influence.

Features	[46]	[86]	[76]	[184]	[18]	[155]	[153]	[70]	[163, 164]	This work
Social network free									✓	✓
Model independent									✓	✓
Non-linear		✓							✓	✓
Temporal dynamics	✓	✓			✓		✓	✓	✓	✓
Indirect influence										✓
Group activity										✓
Platforms of study	Flickr	Plancast, Foursquare	Tencent Weibo	Weibo, CrossFire	Flixster, Twitter, LastFM	Arnetminer, Wikipedia Films	Twitter, Flickr, Arnetminer	Sina Weibo	Twitter	WhatsApp

2.4 Summary

In this chapter, we have discussed prior studies on curiosity modeling and its applications in different areas of Computer Science, notably recommender systems. As discussed in the previous chapters, the existing literature has some important gaps and limitations, some of which we aim at tackling in this dissertation.

Recall that, as argued in the beginning of Section 2.2, the general curiosity appraisal process consists of two key steps, notably, the evaluation of the *stimulation level* based on collative variables and the evaluation of the *curiosity level* based, for instance, on the optimal level of stimulation captured by Wundt’s curve model.

With respect to the first step, the vast majority of computational models exploring psychological theories of curiosity considers only one (or a few) collative variables, thus ignoring various factors that may affect the level of curiosity stimulation on a person as well as the interplay between them. Novelty is by far the most investigated collative variable [186, 143, 179], often captured by metrics related to frequency, recency, dissimilarity and surprise. The latter is an auxiliary variable and, according to Berlyne, does not capture key aspects of novelty [16]. A few other studies focused on different collative variables, notably uncertainty [170] and (social) conflict [179]. In contrast, we here propose metrics to capture all four collative variables proposed by [16], namely, novelty, conflict, uncertainty and complexity, to study online information consumption and production in different case studies.

Moreover, despite not covered by the seminal Berlyne’s work, recent studies have argued for the role of social curiosity as a component of curiosity [120, 80, 52, 63, 61]. Yet, despite the rich literature on social influence analyses and modeling, briefly reviewed in Section 2.3, prior efforts to operationalize its concept as (part of) a collative variable of curiosity stimulation are quite scarce [179, 174, 170], often relying on the explicit social network of users (e.g., friendship relations) and indirectly derived from user ratings. In contrast, we here propose metrics, grounded on information theory, to estimate the importance of social influence as a stimulus to user curiosity driving information sharing. Our metrics are derived for a domain, WhatsApp groups, where no explicit user network exists, and capture both direct and indirect influences that may exist among users on a group.

Concerning the second step of the curiosity appraisal process, namely, the evaluation of the curiosity level based on a given stimulus, prior studies proposed to represent Wundt’s curve of each individual by a bell-shape model, be it a Beta [186] or a Normal [143] distribution, estimating distribution parameters from user behavioral data. Others have modeled the curiosity level as a predictive task, relying on sigmoid functions to model the areas of boredom, anxiety and maximum curiosity. Yet other studies have

opted to skip this step and rather apply the curiosity stimuli, as captured by the collative variables, directly as curiosity scores [174, 170, 6, 111]. By doing so, these studies neglect the possible effects associated with the non-linear relationship between stimulus intensity and curiosity [186]. Also, in general terms, the assumption of a single model of user curiosity may be quite restrictive and unrealistic, given that user curiosity, notably in online applications, can be highly heterogeneous, dynamic and in constant change [83, 80], even considering the same human being. In other words, a person's curiosity may be stimulated quite differently in different situations and may also exhibit long-term changes, reflecting individual evolution. Thus, a broader investigation of the extent to which Wundt's curve is an accurate model of human curiosity, specially for online users, is in great need. In this dissertation, we take a step in this direction by investigating the benefits of considering a combination of distributions to capture the heterogeneity of user curiosity stimulation.

Finally, though the main focus on designing recommendation systems has been on suggesting *relevant* items to the user, a number of recent studies have explored user curiosity scores to better guide the recommendation function towards effective personalized suggestions [186, 1, 143, 179]. Yet, given the limitations of the curiosity computational models adopted by such studies, argued above, these curiosity-driven recommender systems may still have great potential of improvements and can be used in the online information dissemination. In this dissertation, we investigate to which extent the use of multiple collative variables as part of a curiosity stimulation model can offer more fundamental understanding of the information dissemination process.

To that end, we start by presenting our problem statement in the next chapter. Specifically, we introduce the two scenarios covered by our case studies, which relate to complementary perspectives of the information dissemination process: information consumption and information production. We also introduce the main assumptions that guided our work and formally define the problem we tackle.

Chapter 3

Problem Statement and Methodology

As stated in Chapter 1, this dissertation investigates the hypothesis that multi-faceted models of human curiosity, capturing multiple components of curiosity stimulation, can uncover relevant user behavior profiles for the sake of better understanding online information dissemination as well as designing more effective information services (notably personalized recommendation systems). We investigate this hypothesis in light of two complementary perspectives of online information dissemination, notably information consumption and information production. We are particularly interested in studying curiosity as a driving force behind information spread in social media applications, where users play central roles not only in consuming but also in producing and reproducing (i.e., sharing) information.

As such, Figure 3.1 offers an overview of our target problem. Historical traces of user interactions with an online information service (i.e., behavioral data) are used as key input to enable the inference of a user's curiosity at different moments in time. Specifically, we assume that curiosity is a main driving force behind each user interaction with the system, be it the consumption of a particular item (e.g., listening to a song, reading an article) or the production of one (e.g., posting a message). As such, for each new interaction of a user at a given point in time, historical data on previous interactions of the user (e.g., items previously consumed/produced) are used to compute different metrics associated with curiosity stimulation (different collative variables). These metrics can then be used to infer the curiosity stimulus on the user at the interaction time, as well as her expected response to such stimulus (i.e., curiosity level). Such inferences, in turn, can be used to build user curiosity profiles, which can be explored in the design of more effective personalized services.

Driven by the overall scenario depicted in Figure 3.1, we propose to analyze human curiosity in two case studies covering the aforementioned two perspectives of information dissemination. Moreover, based on the general curiosity appraisal process, described in [171], we propose novel metrics capturing different collative variables associated with curiosity stimulation in each such scenario. We also investigate the use of the Wundt's curve as single approach to evaluate human curiosity as a function of the input stimulus.

Before introducing our contributions in the two case studies, which are presented




Figure 3.1: Overview of the problem statement.

in Chapters 4 and 5, we first discuss some general aspects of the problem we intend to tackle in this chapter. We start by introducing in Section 3.1 the scenarios of evaluation corresponding to the two case studies. Next, in Section 3.2, we present some fundamental assumptions on which our user curiosity modeling approach is based. We formally introduce the problem we aim at addressing as well as some general notation in Section 3.3, the individual tasks of the methodology are described in Section 3.4, the key elements of online information services modeling curiosity stimuli are presented in Section 3.5 and briefly summarize this chapter in Section 3.6.

3.1 Case Studies

This dissertation explores the notion that curiosity may propel users to participate in the information dissemination process by consuming previously created content as well as by producing and sharing (new and existing) content. In that case, it seems reasonable that the role of curiosity as such a driving force may vary across platforms, as different application features (e.g., explicit social links, group communication, etc) and content properties may stimulate one's curiosity quite differently [80, 134, 62].

In general terms, the fundamental elements of the information dissemination process that we must be concerned with are:

- *Content items* of information that are produced, consumed or shared. The content items can arouse curiosity in different ways depending on its characteristics, such as the topics of the information content (e.g. different tags of a text message or




Figure 3.2: Main elements of our first case study: online music consumption in LastFM.

document) or any form of categorization of the information content (e.g. music genres, film genres and media types);

- *Users* who interact with the system of online information service by consuming or producing (sharing) content items. Users may act mostly independently or may be influenced by others through implicit or explicit social links;
- *User actions* are the different types of actions a user may have while contributing to information spread (e.g., listening to an audio, watching a video, reading/writing a textual message, sharing an image, etc);
- *Action events* represent specific user actions related to an item's production or consumption. An action event is characterized by a given user producing/consuming a particular content item at a given moment in time.
- *History of user action events*, recorded by the system as logs of user interactions over a given period of time.

Given these elements and the great diversity of possible scenarios, we here choose to focus on two specific case studies, covering platforms with different features, which foster different types of user interactions. Fundamentally the two selected case studies cover different and complementary perspectives of the information dissemination process, namely, information consumption and information production/sharing.

In our first case study, presented in Chapter 4, we investigate the role of user curiosity as a driving force behind users listening online music in one of the largest online




Figure 3.3: Main elements of our second case study: WhatsApp groups.

music service, namely LastFM. In this case, the focus is on information consumption: items are songs characterized by given artists and categories (e.g., different musical genres). The user actions of interest are listening events and user curiosity is assumed to be stimulated primarily by content and individual characteristics, captured in traces of user listenings on the platform. Figure 3.2 illustrates the main elements of this case study.

In our second case study, presented in Chapter 5, we turn our attention to information production and sharing. In that case, we consider a quite different platform, namely WhatsApp (publicly accessible) groups. In such context, the user action of interest is the sharing of content, which can be in different media types (e.g., textual messages, audio, video or image content). Moreover, unlike in the first case study, user curiosity may be stimulated not only by content and individual characteristics but also by the actions of other group members. In other words, we argue that, in a close environment of group communication (as is the case of a WhatsApp group), actions (message sharings) by group members may stimulate others' curiosity and influence them towards particular actions. However, estimating how one's curiosity is stimulated by others' actions is quite a challenge and, in the absence of explicit links connecting group members, must be inferred from prior user actions. Figure 3.3 illustrate the main elements of our second case study.

Having presented our two case studies, we introduce the main assumptions we make to model user curiosity in each such case next.

3.2 Assumptions

Our approach to model user curiosity relies on the derivation of a number of metrics capturing different collative variables associated with curiosity stimulation. These metrics are subsequently used to build user curiosity profiles, which in turn can be applied directly or by means of a function mapping stimuli to curiosity (e.g., the Wundt's curve) as a component of an information service. Towards deriving the proposed metrics, we make three fundamental assumptions that hold for both case studies:

1. The curiosity of a user may be triggered differently depending on various features of the target platform and the types of user actions (e.g., content production, content consumption) under consideration. As such, though grounded on general (information theoretic) guidelines specified by [16] and [144], the derivation of metrics capturing different collative variables should take the particularities of the target scenario into account.
2. The way that the curiosity of a user reacts to a given stimulus may change over time. In other words, a stimulus that triggers great curiosity in a person may later be mostly innocuous, or fall into other regions of the stimulation process (i.e., boredom, anxiety). Thus, the proposed metrics should be (re-) quantified at each action event of a user. Figure 3.2 depicts 9 action events related to item consumption (i.e., song listening in this case). Thus, the level of curiosity stimulation the user is subjected to should be computed at the time of each listening. The same holds for the message sharings shown in Figure 3.3: the level of curiosity stimulus imposed on each user should be recomputed each time the user shares a new message.
3. The curiosity driving an action (be it the consumption, production or sharing of an item) by a given user u at a given time t has a period of activation δ_T , that is, the time interval preceding the user's action event during which previous actions by the user u (and by other users, in case of social influence) may contribute to stimulate u 's curiosity. We refer to that period as a *window of interaction*, and it includes all *action events* that may stimulate u 's curiosity at time t . User actions that occurred before the window of interaction (i.e., prior to time $t - \delta_T$) are assumed to be too old to have any impact on u 's curiosity at time t , being thus disregarded for the sake of computing the curiosity stimulation metrics at that time.

We exemplify the concept of window of interaction in the right side of Figures 3.2 and 3.3. Take the latter for illustration purposes. Assuming $\delta_T = 30$ minutes, the figure shows the window of interaction, notably the messages and users that fall within such window, defined for each message shared by user 3. The messages (and users)

in a window will be used as input to compute the metrics of curiosity stimulation for user 3 at each given time. For example, user 3 shares a textual message (*msg 3*) at 1:20pm. The stimulus to user 3's curiosity driving this particular action will be computed based on the messages shared in the interval [12:50pm; 1:20pm]. Similarly, the stimulus to user 3's curiosity when sharing the audio message at time 2:20pm (*msg 8*) will be (re-) computed based on *msgs 5, 6, 7, and possibly 8* shared within the window [1:50pm; 2:20pm].

In addition to the aforementioned general assumptions, we also rely on a few more premises to design and derive metrics to capture social influence as a collative variable of curiosity stimulation. That is, in our second case study, discussed in Chapter 5, which addresses user curiosity in group communication on WhatsApp, we consider the following additional assumptions:

4. The curiosity of a user may be triggered differently depending on the people with whom the user is interacting and the ongoing discussions among them [120, 62]. In other words, one's curiosity stimulation may vary depending on the particular group the individual is participating in. More yet, if a person participates in multiple groups at the same time, her curiosity stimulation may vary across these groups. Thus our metrics will be computed and our analyses will be performed for different groups separately, considering only the messages shared in each group at a time.
5. The curiosity of a user sharing some content may be stimulated by the other users who shared content in the same group during the window of interaction via *social influence*. We do acknowledge that the content shared itself as well as characteristics of the target user whose curiosity is under analysis may also stimulate one's curiosity. However, when modeling *social* aspects of curiosity, we disregard such characteristics, focusing rather only on all the other *users* who shared content during the window of interaction. Properties of the content and of the target user are exploited to derive metrics related to other collative variables.

As an example, the right side of Figure 3.3 shows the users considered as stimulating user 3's curiosity each time she shares a message. These are all users included in the defined window of interaction *but user 3*. For example, users 1 and 2 are considered to compute the stimulus at time 1:20pm, when user 3 shares her first message. Note that, although user 3 herself is also included in the window of interaction (due to *msg 3* shared at time 1:20pm), only the other users in the window are considered for the sake of analyzing the *social* curiosity driving user 3's behavior. Message 3, shared at time 1:20pm, will be taken into consideration to derive metrics related to other collative variables, as further discussed in Chapter 5. Similarly, only user 2 is considered as stimulating user 3's curiosity when she shares *msg 4* at time 1:40pm.

6. The extent to which user's curiosity is stimulated by other users can be estimated by historical patterns. Specifically, the influence of user u_1 on user u_2 , which relates to the curiosity stimulus triggered on u_2 by u_1 , is proportional to the frequency at which a message shared by u_1 is followed by a message by u_2 within an interval not greater than δ_T (that is, u_1 falls in the window of interaction defined for a message shared by u_2). This frequency is estimated based on the message sharing patterns prior to the current window of interaction (thus, historical patterns).

For example, in the right side of Figure 3.3, the influence of user 1 on user 3 at the time she shares *msg 12* is based on the frequency at which user 3 shared content after user 1 (within a 30 minute maximum interval), *before* the beginning of the current window of interaction, i.e., before 2:30pm. This implies that we will look at all messages shared by user 1 before 2:30pm and check how often user 3 shared a message after it, with a maximum delay of δ_T . In the particular scenario depicted in the figure, it happened twice: *msg 1* at 1:00pm was followed by *msg 3* at 1:20pm, and *msg 6* at 2:10pm was followed by *msg 8* at 2:20pm.

3.3 Problem Definition

We now formally introduce the problem we aim at investigating:

Problem Definition: *Given a set of users $\mathcal{U}$ of an information service and a set of content items $\mathcal{I}$ available on the service, where each item is characterized according to a set of predefined categories $\mathcal{C}$, we aim at quantifying the stimulus a user $u \in \mathcal{U}$ is exposed to when performing an action¹ on a content item $i \in \mathcal{I}$, characterized by (one or more) categories $c \in \mathcal{C}$, which serve as proxy for representing content properties.*

To that end, we we will explore a set of chronologically ordered tuples $(u, i, \mathcal{C}_s, t)$, where each tuple specifies that a user u acted on item i , characterized by the categories $\mathcal{C}_s \subseteq \mathcal{C}$ at time t . In other words, the set of tuples build a *history of users' action events* from which user curiosity (stimulation) will be inferred at each action time t . We assume $|\mathcal{U}| = n_{users}$, $|\mathcal{I}| = n_{items}$ and $|\mathcal{C}| = n_{cat}$.

¹We note that, since in both case studies, user actions are of a single type, either music listenings (case study 1) or message sharings (case study 2), we do not consider the action type as a dimension of exploration and, as such, do not explicitly consider it in our notation. Yet, future extensions may consider different action types in a single scenario of study. In that case, notation must be adjusted to include them elements of the problem definition.




Figure 3.4: Methodological areas and tasks of this dissertation: (a) *blue blocks* represent the tasks and (b) *gray documents* represent the artifacts (i.e., outputs) of the respective tasks.

3.4 Methodology

Figure 3.4 shows the methodology blocks and tasks used to investigate our two case studies. Furthermore, this methodology can be applied to any type of online information service. This methodology consists of three different areas, namely (i) *evaluation of stimulus level*, (ii) *evaluation of curiosity level* and (iii) *evaluation of model validation*. The first two areas of Figure 3.4 are directly related to the *general appraisal process* (see Section 2.2), i.e., the first area deals with the evaluation of the stimulus level based on the stimulus captured by the collative variables (e.g., traditional collative variables or social variables). The second area, in turn, is related to the evaluation of the curiosity level, which deals with the optimal level of stimulation, i.e., Wundt's curve (the curiosity profile). And the last area of this figure deals with the validation of our proposal for a complete curiosity model.

We begin by explaining the *Evaluation of the stimulus area* in Figure 3.4. To this end, the first task is the *Instantiation of metrics* and it is important to define the metrics for the traditional collative variables (i.e., novelty, uncertainty, conflict and complexity) and for the social curiosity variables. In this dissertation, when presenting the case studies, we will define how the stimulus can be captured by these metrics. Since these metrics are defined, we need to instantiate them for each event of the user action, be it content production or content consumption. In this context, instantiation means that depending on the characteristics of the information content accessed by the user, a specific instance

of the metric is created and thus the stimulus value for that access is captured from the user's action event. Next, the task of *Selection of metrics* aims to find out which metrics capture complementary aspects of the stimulus of curiosity. For this purpose, the Spearman correlation coefficient is calculated for each pair of metrics from each user. It is then important to determine the percentage of users whose correlation coefficient lies in the interval $(-1.0, -0.5)$ and $(+0.5, +1.0)$. In this dissertation, we consider the correlation coefficient within this interval as high or moderate correlation. Then, the pairs of metrics that have more than 70% of cases with high/moderate correlation are considered as redundant pairs and only one metric is selected from these pairs. Therefore, only the metrics that are considered complementary should be used in the next task. Finally, in these figure the *complementary metrics* are presented as an artifact of the *Selection of the metrics* task.

After the selection of metrics, in Figure 3.4, there are three tasks related to the characterization of the stimulus, divided into three levels, namely:

- (1) *Characterization at the access level*, this task uses only the selected *complementary metrics* from the previous task and aims to identify different profiles of stimulus patterns through clustering techniques. The outputs (or artifacts) of this task are therefore the clusters (or profiles) of stimulus patterns (this is represented by the gray document below this task in Figure 3.4);
- (2) *Characterization at the individual level* which includes the characterization of users based on the stimulus pattern profiles from the access level. This means that each user is represented by a set of states based on the stimulus pattern profiles that show how each user switches between different profiles. From this set of states, clustering techniques should therefore be used to identify the *profiles of the users* as a result of this task (see the gray document below for this task in Figure 3.4); and
- (3) *Characterization at the collective level*, this involves only the social influence metrics considered to the social curiosity. We emphasize that this task can only be realized if the *social network* exists with *social connections of friendship* implicitly or explicitly present in the application environment (i.e., the online information service). The former, there are *implicit social connections* in group communication platforms such as WhatsApp (our second case study). In this case, interactions such as the sharing of messages by users in a group implicitly presuppose the existence of a social network. The latter, *explicit social connections* are of course present in social media such as Instagram or X/Twitter, these social connections are intrinsic to these social media. For this task, it is therefore necessary to aggregate the individual metrics of the user, which are measured by each of his action events (content consumption or content production). In this way, a measure can be calculated that represents user behavior at a collective level and can be used to identify collective behavior

(as will be seen in the second case study in Chapter 5). Consequently, the artifact of this task is the *Collective stimulus profile* given by an aggregate metric of social influence (the gray document below of this task in Figure 3.4).

We emphasise that these three tasks of characterization are independent of each other. Therefore, they can be performed in isolation or one after the other. Figure 3.4 shows a diamond shape between these three tasks of characterization, indicating two possible parallel output streams. The task of *Characterization at access level*, for example, is followed by the two parallel output streams, the task of *Characterization at the individual level* and the task of *Modeling the curiosity profile*. The latter belongs to the second area of the methodology – which corresponds to the middle of Figure 3.4, the *Curiosity level evaluation* (it is related to the second step of the general appraisal process).

To continue the explanation, the task of *Modeling the curiosity profile* in the *Curiosity level evaluation* area in the middle of Figure 3.4 is to estimate the Wundt’s curve that represents the user’s curiosity profile from all the stimuli measured for each user action event, i.e., from the user’s access history. This idea is inspired by [186] (and described in Section 2.2.2.2 of Chapter 2), where the stimulus value measured for each user access by the metric of collative variables is used to estimate the Wundt’s curve – this includes the traditional collative variables and the social curiosity variables. We recall that each access of a user has different aspects of curiosity stimulus, which are captured by the set of metrics. Then, all of the user’s stimuli obtained from user access history are entered into a probability distribution function that maps the stimulus to a curiosity score according to its frequency. A Beta mixture distribution is used for this purpose. This is our main difference with [186], because these authors assume that each user can be modelled by a single Wundt’s curve. To this end, the average of the different aspects of the user access stimulus must be determined, i.e., the average of the different metrics used to capture the different aspects of the stimulus. This average is then mapped to the range $[0, 1]$ via a min-max normalization, since the Beta distribution has the domain in this interval. Next, the set of averages of stimuli from all user accesses is used to estimate the curve of the user’s curiosity profile, i.e., Wundt’s curve. More details on this task are described in the 4 and 5 chapters, where we look at our case studies. The generated artifact is, of course, the *user’s curiosity profile* represented by the gray document associated with the task of *Modeling the curiosity profile* in Figure 3.4.

The next task in Figure 3.4 is the *Analysis of curiosity profile and stimulus pattern analysis*, which deals with analyzing the relationship between the user’s curiosity profile and the stimulus pattern profile. This task is about understanding how the classes of curiosity profiles (i.e., Wundt’s curve) can be explained by the composition of different types of stimulus pattern profiles. For this purpose, one must first know that there are three possible classes of users’ curiosity profiles, namely: (i) a single Wundt’s curve estimated by a Beta distribution; (ii) multiple Wundt’s curves estimated by a 2-Beta

mixture (multivariate) distribution; and (iii) multiple Wundt's curves estimated by a 3-Beta mixture (multivariate) distribution. The first step is therefore to calculate in aggregated form the proportion of each stimulus pattern profile for each of the three classes of Wundt's curve. This step aims to verify how these classes are formed by different stimulus patterns. The second step, on the other hand, proposes to break down each of the user's curiosity profiles individually and showing how the curve of the curiosity profile is shaped by the different curves of the individual stimulus pattern profiles. This can show which stimulus patterns dominate the formation of the multiple Wundt's curve and how they influence the curve formation for each user. The last step focuses only on the users modeled by the multiple Wundt's curve, i.e., the 2- and 3-Beta mixture distribution. With this in mind, the distribution of the proportion of accesses per stimulus pattern profile should be analyzed when normalized per user. This is to determine if there are some stimulus pattern profiles that increase when users are modeled by multiple Wundt's curves – these users represent more complex behaviors.

The penultimate task is the *Temporal analysis of the curiosity profile* in the *Curiosity level evaluation* area of Figure 3.4, which focuses on users modeled by multiple Wundt's curves. This task aims to check whether users are still modeled by a Beta mixture distribution when their entire activity period (i.e., historical of accesses) is divided into short time periods. To do this, a few users should be selected for this analysis. Then the entire period of each user's access history is divided into short periods such as weeks. So, for each short period, the Wundt's curve is estimated based on the accesses within these short periods. Finally, for each estimated Wundt's curve, it is checked whether it was estimated using a 2- or 3-Beta mixture distribution. If the majority of Wundt's curve is indeed estimated by Beta mixture distributions, then this shows that user behaviour remains complex regardless of the size of the time period analysed.

We conclude this section with the last task, which belongs to the third area related to the *Model validation evaluation*. The evaluation is called *Validation of the curiosity model* in Figure 3.4 and should show that the metric does not capture random stimulus values, i.e., that they originate from real user behavior and did not arise by chance. This task therefore involves the definition of hypotheses that are to be tested with the help of null models. It is important that these hypotheses are based on the fundamental assumptions mentioned in Section 3.2. Therefore, a small sample of users with representative activity is selected and the null models should consist of random changes in the users' access history that aim to violate the defined hypothesis. Finally, the random stimulus values generated by different null models should be compared with the real stimulus values of the users. The greater the difference between these values, the better the result. This shows that the real stimulus values are not caused by chance.




Figure 3.5: Structure for modeling stimulus of curiosity in social media with its key elements: (a) *user* as central element; (b) *content items* of information; (c) *social connections* in the system; (d) *user's actions* as different types of actions including social interactions; (e) *history of user accesses* recorded by the system.

3.5 Fundamental structure and key elements

The purpose of this section is to show how we can apply stimulus curiosity modeling to any type of online information service, especially social media. To this end, it is important to define the key elements of the system such as (i) information content and (ii) resources provided by the system that users can consume, product and interact, (iii) the main characters (i.e., users) and (iv) their respective actions.

Figure 3.5 show the fundamental structure for modeling curiosity stimuli that is composed by the key elements of social media where the *user* is highlighted as central element who interact with the system of social media whose accesses the *content items* of information available as resource of system with different characteristics such as categories. Moreover, this user has *social connections* that are established in the system with their respective social influence, which depends on social capital. Yet, there are *user's actions* as different types of actions which contribute to information spread such as consumption of content items, production of content items and interactions with others through implicit or explicit social links. At end, the *history of user accesses* are recorded by the system with data of accesses such as user identification, content item identification, categories of content item and timestamp.

The first step in modeling is to understand the possible resources available in the systems. Regardless of the type of system, it has *content item* of information – they are shown in Figure 3.5, top left. In this sense, it is important to determine the *categories* of the content items. These categories are very important because the metrics based on

the traditional collative variables depend on them to capture the different amounts of information available in each category beyond the content item itself. Intuitively, the categories represent the *diversity* of information in a content item. Another important resource in the system is what the user can do with the content item. We call this the *user action*, which can be the information *consumption* or information *production* of content items (these user actions are shown to the right of the user in the center of Figure 3.5). In this dissertation, we define information consumption as the action that the user performs with the content item, e.g., the way the user uses or processes the content that has aroused his or her curiosity. We also refer to this type of information consumption as *explicit*. On the other hand, *implicit* information consumption can be a reaction to the information content, such as a like or dislike, which is also a resource of the system. Information production, in turn, is the action that the user performs to produce or create a content item in the system. To do this, the user must use the available system resources, e.g., the various media file types supported by the system, such as video, audio, image or text, to name but a few.

Another type of user action that can occur in the system, especially in social media, are *social interactions* (this is shown to the right of the user in the center of Figure 3.5). These actions are naturally performed in social media where it is possible to create social networks. The social connections of friendship (they are shown in Figure 3.5, bottom left) can be *implicit* or *explicit* depending on the social media. As we have seen in the previous section on collective-level characterization, an important task of the methodology, explicit social connections emerge when there are friendships between users and they perform social interactions through the available resources of the social media, such as creating a post of a media type or sharing a message. Implicit social connections, on the other hand, usually occur when users do not have a clear friendship, but have social interactions via social media resources such as sharing a message or replying to a message in the system. Examples of such systems are group communication platforms (e.g., WhatsApp or Telegram) and online networking communities (e.g., Reddit communities and forums² or Discord community server³) where users can share common interests.

Furthermore, we emphasize that all these user actions (i.e., information consumption, information production, and social interactions) are recorded by the system as logs of user interactions over a period of time. These logs are therefore referred to as the history of user accesses and are shown in the right side of Figure 3.5.

Finally, it is important to note that depending on the system, assumptions as defined in Section 3.2 may need to be adjusted or some assumptions should be created according to the characteristics and resources of the online information services.

²<https://support.reddithelp.com/hc/en-us/articles/204533569-What-are-communities-or-subreddits>

³<https://support.discord.com/hc/en-us/articles/14078261239831-Clarifying-Server-Types>

3.6 Summary

In this chapter, we have discussed general elements of the work developed in this dissertation. We have introduced the two case studies as well as some fundamental assumptions that we use to guide our developments. We also give a formal definition of our target problem and a complete description of our methodology. We conclude the chapter with a description of modeling the curiosity stimulus in online information services. All these elements serve as pillars for the information offered in the next two chapters where we present the studies performed on each scenario of investigation. We start our discussion in Chapter 4 with our first case study, focused on information consumption. Specifically, we present our efforts to model user curiosity as a force behind user's listening to online music.

Chapter 4

Analyzing and Modeling User Curiosity in Online Content Consumption: A LastFM Case Study

This chapter is organized as follows. The next section presents the introduction of our first case study. The Section 4.2 discusses the modeling of user curiosity. Next, we present our results addressing the aforementioned research questions in Sections 4.3 and 4.4. Our methodology and the experiments to validate the metrics by using null models are described in Section 4.5. Finally, Section 4.6 provides a discussion of the results, shows a limitation of work and possible directions for future work.

4.1 Introduction

Our goal here is to model and analyze user curiosity as a driving force behind user consumption of online content. Building upon existing models [186], we consider all collative variables as sources of stimulus to curiosity. Also, we investigate whether a single Wundt’s curve is a reasonably model of user curiosity. As case study, we focus on the consumption of online music from LastFM, as musical tastes are related to personality characteristics [48]. To that end, we analyze a large dataset consisting of over 243,555,074 listening events, covering more than 84,466 users and 9,319,305 songs. Specifically, we address two research questions:

- **Research Question 1 (RQ1):** Are there distinct user behavior profiles in terms of curiosity stimuli, as captured by the collative variables?
- **Research Question 3 (RQ3):** To which extent user curiosity can be accurately modeled by a Wundt’s curve?




Figure 4.1: Levels of content categorization for song (content items).

Given our focus on online music, we start by proposing stimulus metrics related to songs, artists and their musical genres capturing the collative variables. Next, we find that, though the single Wundt's curve is a reasonable model, for more than 40% of them the curiosity curve is multi-modal, reflecting a mixed behavior. So, it can be well modelled by the combination of two or three Wundt's curves.

In sum, this chapter presents a thorough effort of modeling user curiosity as a driving force of online music consumption. Compared to prior studies, our contributions are: (1) the proposition and analysis of a set of metrics that capture different perspectives of the stimuli offered to users as they are exposed to online content; and (2) new models of curiosity that more accurately account for the diversity of user behavior.

4.2 Modeling User Curiosity

In this section, we present our notation used to define our metrics and, then, we present metrics for each one of collative variables of *novelty*, *uncertainty*, *complexity* and *conflict* in the context of LastFM.

4.2.1 Notation

Our context of analysis is composed by sequences of songs consumed by users in LastFM dataset, a platform of online music consumption. So, a certain song belongs to

Table 4.1: Basic notation used to derive the curiosity stimulation metrics.

Notation	Description
$\mathcal{U}$	set of all users (individual user indicated by u)
$\mathcal{S}$	set of all songs (one group indicated by s)
$\mathcal{A}$	set of all artists (one group indicated by a)
$\mathcal{C}$	set of all musical genres (i.e., categories) used as classification of artists or band
$\mathcal{C}_a$	set of musical genres associated with a given artist of listened song (one category is indicated by c)
$t_{i u}$	timestamp of the i^{th} song consumed by user u
δ_T	duration of window of interaction
u	user whose curiosity we are analyzing
$\rightarrow$	indicates some quantity computed for current window of interaction
$\mathcal{S}_t^{\rightarrow}$	set of (distinct) songs consumed during window $[t - \delta_T, t]$
$\mathcal{A}_t^{\rightarrow}$	set of (distinct) artists from songs consumed during window $[t - \delta_T, t]$
$\mathcal{C}_t^{\rightarrow}$	set of (distinct) musical genres of artists from songs consumed during window $[t - \delta_T, t]$
$\mathcal{W}_{\delta_T}$	set of listening events with distinct $t_{i u}$ that occur within during window $[t - \delta_T, t]$
S	random variable associated with songs
A	random variable associated with artists
C	random variable associated with musical genres
$n_{s t}^{\rightarrow}$	number of times that song s was consumed during window $[t - \delta_T, t]$
$n_{a t}^{\rightarrow}$	number of times that artist a occurs for consumed songs during window $[t - \delta_T, t]$
$n_{c t}^{\rightarrow}$	number of times that musical genre c occurs for consumed songs during window $[t - \delta_T, t]$

a certain artist, which is considered to have a *content categorization*, which in turn has a *categorization* of certain musical genres. This categorization of songs in terms of artists and music genres is defined as *level of categorization* and can be visualized in Figure 4.1. In that picture, one can see a possible structure of these categorization levels in form of a tree, where: (a) the root node is the song; its respective child node is (b) the artist (i.e., the first level of content category); and, next below, there are (c) k leaves, child nodes of artist, which denote the musical genres, it denotes the second level of content category. We emphasize that an artist may have one or more musical genres.

Our present goal is to analyze and model user curiosity in the particular context of online music consumption. To that end, we consider the temporal sequence of listening events of a user u . Each listening event of user u is associated with a timestamp t and a unique song (track) s by an artist (singer/band) a . We use the existing musical genres of LastFM to categorizing the artists, then we define $\mathcal{C} = \{\text{"rnb"}, \text{"rap"}, \text{"eletronic"}, \text{"rock"}, \text{"new age"}, \text{"classical"}, \text{"reggae"}, \text{"blues"}, \text{"country"}, \text{"world"}, \text{"folk"}, \text{"easy listening"}, \text{"jazz"}, \text{"vocal"}, \text{"children's"}, \text{"punk"}, \text{"alternative"}, \text{"spoken word"}, \text{"pop"}, \text{"heavy metal"}\}$.

The set $\mathcal{C}_a$ includes all musical genres used to classify the artist/band of a song. Thereby, each artist a has one or more tracks and is characterized by k musical genres with $1 \leq k \leq n_{cat}$, wherein $|\mathcal{C}| = n_{cat}$, and, yet, $\mathcal{C}_a \subseteq \mathcal{C}$. In turn, the set $\mathcal{U}$ denotes the set of users whom consume songs in LastFM platform and $u \in \mathcal{U}$. Thus, each listening event is defined by a tuple $(u, s, a, \mathcal{C}_a, t)$.

Also, the previous definitions demonstrate the basis of our notation. Variables are

represented as small letters (e.g., user u and song s), random variables as capital letters (e.g., S and A to song and artist, respectively), and sets are represented as calligraphic letters (e.g., $\mathcal{S}$ and $\mathcal{A}$, song and artist, respectively). We employ subscripts to determine subsets (e.g., songs of specific user u as S_u). Finally, we use the Bayesian notation to define constraints/dependencies (e.g., $t_{i|u}$ is the time of the i -th listening event from a given user u). The complete set of notation, including those used above as others introduced in the following sections, is listed in Table 4.1.

We deal with the final metrics in a hierarchy to show how they relate to one another: songs, artist and genres. This hierarchical approach is crucial to understand exactly which concepts drive user curiosity. That is, some users may be curious towards more specific concepts, such as songs, while others focus on general concepts such as genres. For genres, the metrics complement each other by capturing different aspects of the same concept, as we will see in the next section.

4.2.2 Computing Stimuli

Our proposal of metrics for collative variables follows the guidelines of [16] and is therefore based on measures of information theory. In this sense, the measures of information theory can be applied when an event is selected from an *information space* that is a set of alternative, mutually exclusive events with a probability for each of them. According to [16], events that are selected from two *information spaces* are referred to as *input signals* and *output signals* and information theory measures can be applied. In this case, these measures relate to the transmission of information and reflect the degree of correspondence between the two information spaces. From this perspective, there are two *information spaces* in psychology that can be treated as *input* and *output* of communication channel [16]. The former is the *stimulus space* and the latter is the *response space*. The *stimulus space* therefore refers to the set of possible stimuli that a person can receive. The *response space*, on the other hand, refers to the set of possible responses that a person can show in response to a selected stimulus.

In this chapter, our goal is to analyze and model user curiosity in the specific context of online music consumption. Therefore, the *stimulus space* in an online music consumption system such as LastFM can be denoted by the set of songs a user can listen to. Otherwise, the *response space* can be represented by the set of songs that the user has actually consumed. In this context, a *response* can therefore be seen as the act of listening to a song.

In psychology, there are many variables that are of interest to the psychologist,

such as how quickly someone responds to a stimulus, how accurately someone identifies a stimulus, or how likely it is that someone will remember a stimulus correctly after it has disappeared [16]. These variables are influenced by the type of stimulus that occurred and what stimuli/responses could have occurred instead and how likely they were. This means that they depend on the degree of *uncertainty*.

In order to calculate the degree of uncertainty, the information spaces must be described. For this purpose, as defined in Section 2.1, (i) a list of stimuli that can occur should be created, (ii) these should be divided into classes and (iii) a probability should be defined for each class. In this sense, the probability of each class indicates the competing strength of the possible responses and the degree of uncertainty is given by the Shannon *entropy*, i.e., the average amount of information of *stimulus space* [92].

Yet, there is a measure of information theory that indicates the information content of a stimulus, namely the Shannon *information content*, also known as surprisal and is a natural measure of information content [92]. This measure depends on the probability or relative frequency of the class to which the stimulus belongs. The lower this probability, the less information the stimulus contains. The *information content* of a stimulus of a certain class is defined as the negative logarithm to the base 2 (i.e., $-\log_2$) for the probability of this class. This measure is regarded as a measure of *novelty*. All our metrics are defined in the unit *bits*.¹ This was a design decision to facilitate the aggregation of metrics.

In this context, we begin by stating that one important factor when understanding user behavior online is time. In fact, temporal aspects such as recency (e.g., last time an item was accessed) will naturally correlate with our variables. In this sense, all metrics are computed for a given user u and specific listening event i , using as input all previous listening events in the window $[t_{i|u} - \delta_T, t_{i|u}]$. Thus, the curiosity of the user for access i depends on every previous access inside the window, i.e., the *stimulus space*.

Based on such temporal aspect from assumption 3 (see Chapter 3), it is important to consider that users' curiosity will likely dwindle after long periods. In our analysis, we therefore consider a $\delta_T = 24$ hours. That is, the curiosity of the user depends **only** on the events from the last 24 hours, that is, during the window of interaction (thus the use of notation $\rightarrow$ in this section). This choice is based on the seminal work of [186], in which the authors work in the same context of online music consumption. We leave the evaluation of other size windows as future work.

We define six metrics based on Shannon *information content* as follows. With s , which captures the current song (track), a the current artist and $\mathcal{C}_a$ the current musical genres, we define $songNovelty(u, s, t = t_{i|u})$, $artistNovelty(u, a, t = t_{i|u})$ and

¹*Bit* is the smallest unit of information in a computer that can be stored or transmitted, and is denoted by a binary variable whose value is 0 or 1, *bit* in turn is also used as a unit of information in information theory.

$genreNovelty(a, \mathcal{C}_a, t = t_{i|u})$ as the novelty of tracks, artists and musical genres for a particular artist. The novelty of given song is defined as:

$$songNovelty(u, s, t = t_{i|u}) = \begin{cases} -\log_2(P_t^\rightarrow(S = s)), & \text{if } P_t^\rightarrow(S = s) > 0 \\ -\log_2(1/|\mathcal{W}_{\delta_T}|), & \text{otherwise} \end{cases} \quad (4.1)$$

where $P_t^\rightarrow(S = s)$ is the probability of song s being consumed during the window of interaction, which is defined as $n_{s|t,u}^\rightarrow / \sum_{s \in \mathcal{S}_u} n_{s|t,u}^\rightarrow$, where $n_{s|t,u}^\rightarrow$ denotes the number of times which the song s appeared into window of interaction of user u at time t . The smallest value for a metric based on the Shannon information content is 0 (zero) bits, i.e., no information, if $P_t^\rightarrow(S = s) = 1$. Or if $P_t^\rightarrow(S = s) = 0$, the metric is set to the largest surprisal associated with a current song s (or another current *content item* such as the current artist a or the current musical genres $\mathcal{C}_a$). In this case, we set the metric to the uniform distribution of an *action event* within the window of interaction $[t - \delta_T, t]$, i.e., $P_t^\rightarrow(S = s) = 1/|\mathcal{W}_{\delta_T}|$, where $\mathcal{W}_{\delta_T}$ is the set of *action events* (i.e., listening events) with distinct $t_{i|u}$ that occur within this window. This means that $1/|\mathcal{W}_{\delta_T}|$ is the smallest probability for an *action event* within the window. Therefore, we set the largest value for all our metrics based on the Shannon *information content*, i.e., $-\log_2(1/|\mathcal{W}_{\delta_T}|) > 0$.²

Similarly, artist novelty is defined by looking at artists instead of songs. Both metrics captures will be higher for less accessed artists or songs. So, it is given by

$$artistNovelty(u, a, t = t_{i|u}) = \begin{cases} -\log_2(P_t^\rightarrow(A = a)), & \text{if } P_t^\rightarrow(A = a) > 0 \\ -\log_2(1/|\mathcal{W}_{\delta_T}|), & \text{otherwise} \end{cases} \quad (4.2)$$

where $P_t^\rightarrow(A = a)$ is the probability of artist a to appear for given song consumed into window of interaction. These probability is defined as $n_{a|t,u}^\rightarrow / \sum_{a \in \mathcal{A}_u} n_{a|t,u}^\rightarrow$, where $n_{a|t,u}^\rightarrow$ denotes the number of times which the artist a appeared into window of interaction of user u at time t .

Given a particular artist of a consumed song at time t , let C be the random variable associated with musical genre, c be a musical genre associated with the given artist from a song (i.e., $c \in \mathcal{C}_a$). We recall that $\mathcal{C}$ contains all possible existing musical genres into LastFM platform, in total, there are 20 musical genres (see Section 4.2.1). In the same way, as defined for previous metrics, we states the novelty metric related to musical genres,

²We claim that our metric is not the Shannon *information content* in its original form, but it is based on it. The reason for this is that the sum of the probabilities for all content items in a window should be 1, and we use the probability of a content item to calculate Shannon *information content*. However, if the probability of a content item is zero, we set the metric to $-\log_2(1/|\mathcal{W}_{\delta_T}|)$, where the probability of $1/|\mathcal{W}_{\delta_T}|$ is not considered when summing the total probability for all content items within the window. In this way, we ensure that no content item that is in the window can have a higher Surprisal value than one that is not in the window.

grounded on the surprisal associated with that variable:

$$\text{genreNovelty}(a, \mathcal{C}_a, t = t_{i|u}) = \begin{cases} -\log_2 \left(\bar{P}_t^\rightarrow(\mathcal{C}_a) \right), & \text{if } \bar{P}_t^\rightarrow(\mathcal{C}_a) > 0 \\ -\log_2 (1/|\mathcal{W}_{\delta_T}|), & \text{otherwise} \end{cases} \quad (4.3)$$

where $\bar{P}_t^\rightarrow(\mathcal{C}_a)$ is the average probability of musical genres associated with the artist (musical genres in $\mathcal{C}_a$), and it is given by

$$\bar{P}_t^\rightarrow(\mathcal{C}_a) = \frac{1}{|\mathcal{C}_a|} \sum_{c \in \mathcal{C}_a} P_t^\rightarrow(C = c),$$

yet, the probability of a specific musical genre is estimated as $P_t^\rightarrow(C=c) = n_{c|t,a}^\rightarrow / \sum_{k \in \mathcal{C}_t^\rightarrow} n_{k|t,a}^\rightarrow$, wherein $n_{c|t,a}^\rightarrow$ is the number of musical genres equal to c of an artist a from a song consumed in the window of interaction $[t - \delta_T, t]$. That metric captures how surprising it is that the song consumed at time $t = t_{i|u}$ has musical genre c .

In this work, we use one of the four criteria, see the list of these criteria in the Section 2.1, proposed by Berlyne to measure conflict, namely the absolute strength of the competing responses [16]. This strength is represented by the probability of a class. So, we propose to capture the conflict associated with given listening event based on the probabilities of the user listening to different musical genres (absolute strength), i.e., $P_t^\rightarrow(C=c)$. More specifically, the conflict is estimated as follows:

$$\text{genreConflict}(a, \mathcal{C}_t^\rightarrow, t = t_{i|u}) = -\log_2 \left(\bar{P}_t^\rightarrow(\mathcal{C}_t^\rightarrow) \right), \quad (4.4)$$

where $\bar{P}_t^\rightarrow(\mathcal{C}_t^\rightarrow)$ is the average probability of music genres calculated over all genres that the user has listened to in the interaction window defined by

$$\bar{P}_t^\rightarrow(\mathcal{C}_t^\rightarrow) = \frac{1}{|\mathcal{C}_t^\rightarrow|} \sum_{c \in \mathcal{C}_t^\rightarrow} P_t^\rightarrow(C = c). \quad (4.5)$$

In the same way, we define the two complexity metrics, instantaneous and overall complexity, denoted by *instGenComplex* and *overallGenComplex* respectively, as:

$$\text{instGenComplex}(\mathcal{C}_a, \mathcal{C}_t^\rightarrow, t = t_{i|u}) = -\log_2 \left(\frac{|\mathcal{C}_a|}{|\mathcal{C}_t^\rightarrow|} \right) \quad (4.6)$$

and

$$\text{overallGenComplex}(\mathcal{C}_t^\rightarrow, \mathcal{C}, t = t_{i|u}) = -\log_2 \left(\frac{|\mathcal{C}_t^\rightarrow|}{|\mathcal{C}|} \right), \quad (4.7)$$

where $|\mathcal{C}|$ is the total number of distinct genres in the complete LastFM platform. The instantaneous complexity is simply the fraction of the number of genres in the current song, $|\mathcal{C}_a|$, divided by the number of genres consumed so far (including this song), $|\mathcal{C}_t^\rightarrow|$. The overall complexity is measured as fraction of the total number of distinct genres listened so far, i.e. $|\mathcal{C}_t^\rightarrow|$, that exist inside window of interaction divided by the $|\mathcal{C}|$. The

instantaneous complexity is related to the current song. It answers question's such as: how complex is this song when compared to the other songs consumed in the window of interaction. In contrast, the overall complexity is related to the window of interaction and answers questions as: How complex is the behavior of the user in the last 24 hours.

As follows, there are properties of metrics for collative variables that are based on the Shannon *information content*:

Properties of metrics based on surprisal: (1) the lower this probability (i.e., competing strength) is, the more information the stimulus contains, and (2) no stimulus within the interaction window can have a higher surprisal than a stimulus that is not within the window of interaction.

The first property of the six surprisal-based metrics states that the closer the probability is to zero, i.e. the more information they capture, the greater the value of these metrics. Otherwise, the closer the probability of a class is to 1, the lower the value of these metrics, i.e., the less information they capture.

The second property in turn belongs to the novelty metrics (i.e., novelty of song from Equation 4.1, novelty of artist from Equation 4.2 and novelty of genre from Equation 4.3) and is considered for the case that the current song is not within the window of interaction, that is, the surprisal is the greatest. In other words, this is denoted by the “otherwise” part of the equations of novelty.

In addition, the properties of all metrics based on surprisal are explored by axioms which are defined in the Appendix A in the Section A.1 where they are exemplified for better understanding.

Finally, it should be noted that unlike conflict, which capture the diversity of the song in terms of the artist’s musical genres when only those included in the window of interaction are considered, complexity again captures a slightly different notion of diversity, which considers (i) all possible musical genres in the current song and (ii) all possible genres listened in the window in terms of the entire genres of LastFM.

To capture the uncertainty, we rely on the probabilities that the user listens to different musical genres within the window of interaction $[t - \delta_T, t]$. The probability of a class (i.e., musical genres), i.e., $P_t^\rightarrow(C = c)$, denotes the strength of the corresponding response, whereby the uncertainty of a stimulus refers to the Shannon entropy of the classes. Specifically, uncertainty is quantified by the entropy of musical genres. Let $P_t^\rightarrow(C = c)$ as previously defined, it is calculated for all genres that the user has listen in the window of interaction (i.e., genres in the set $\mathcal{C}_t^\rightarrow$). The uncertainty in the current listening event is defined as:

$$\text{genreUncertainty}(\mathcal{C}_t^\rightarrow, t = t_{i|u}) = -\sum_{c \in \mathcal{C}_t^\rightarrow} P_t^\rightarrow(C = c) \log_2 \left(P_t^\rightarrow(C = c) \right). \quad (4.8)$$

Table 4.2: Metrics of curiosity stimulation for music consumption in LastFM.

Collative Variable	Metric	Definition
Novelty	$songNovelty(u, s, t)$	Novelty associated with user u whose consumed a given song s at time t (Equation 4.1)
Novelty	$artistNovelty(u, a, t)$	Novelty associated with user u whose consumed a given song of an artist a at time t (Equation 4.2)
Novelty	$genreNovelty(\mathcal{C}_a, t)$	Novelty associated with set $\mathcal{C}_a$ of musical genres of an artist a with respect to the song consumed by user u at time t (Equation 4.3)
Conflict	$genreConflict(\mathcal{C}_t^{\rightarrow}, t)$	Conflict associated with set $\mathcal{C}_t^{\rightarrow}$ calculated over all genres c with respect to the all songs consumed in the window of interaction (Equation 4.4)
Complexity	$instGenComplex(\mathcal{C}_a, \mathcal{C}_t^{\rightarrow}, t)$	Instantaneous complexity associated with set $\mathcal{C}_a$ of musical genres of an artist a with respect to the song consumed by user u at time t (Equation 4.6)
Complexity	$overallGenComplex(\mathcal{C}_t^{\rightarrow}, \mathcal{C}, t)$	Overall complexity associated with set $\mathcal{C}_t^{\rightarrow}$ of musical genres from all consumed songs in the window of interaction measured at time t (Equation 4.7)
Uncertainty	$genreUncertainty(\mathcal{C}_t^{\rightarrow}, t)$	Uncertainty associated with set $\mathcal{C}_t^{\rightarrow}$ of musical genres from all consumed songs in the window of interaction measured at time t (Equation 4.8)

As with the metrics based on *surprisal*, we also define some properties of the metric based on Shannon *entropy*:

Properties of metric based on entropy: (1) *this measure reaches the maximum when all probabilities (i.e., competing strength) are equal, and (2) when they are equal, it increases with the number of alternative classes to which the stimulus could belong.*

These two properties mean that the probabilities of the classes in the window of interaction are equal, and the more classes there are, the greater the uncertainty of the musical genres. Moreover, these properties are explored through axioms defined in the Appendix A, more precisely in the Section A.1, where they are explained in detail.

In short, the Table 4.2 summarizes the metrics for stimulating curiosity about online music consumption on the LastFM platform. We emphasize that these metrics should be interpreted as: *the larger the value, the greater the curiosity stimulation.*

Table 4.3: Statistics of LFM-1B Dataset.

Item	Number
Users	84,466
Artists	219,055
Albums	4,351,218
Tracks	9,319,305
Listening events	243,555,074
Users with at least 1K listening	43,220

4.2.3 The LastFM Dataset

As a case study, we employ the LFM-1B Dataset of [131]. This dataset captures the listening behavior of users on the LastFM website. The data covers the period of January 2013 to August 2014. Each event recorded contains the following features: user, artist, album, track and timestamp of listening. To capture complexity, we incorporate into the dataset the genre of artists (e.g., The Beatles belongs to the rock genre). In particular, we make use of the genres defined in [132]. It is comprised of: *rnb*, *rap*, *eletronic*, *rock*, *new age*, *classical*, *reggae*, *blues*, *country*, *world*, *folk*, *easy listening*, *jazz*, *vocal*, *children's*, *punk*, *alternative*, *spoken word*, *pop* and *heavy metal*. Table 4.3 summarizes our dataset.

Recall that we compute the metrics over time windows of 24 hours. In order to avoid data sparsity issues, we consider in our analysis only users with at least 1,000 listening events in the whole period. We also only consider events with at least 30 previous listening events in each 24-hour window between January 2013 to August 2014. After such filtering, we were left with 43,220 users with 149,949,511 stimulus events.

4.3 Understanding Metrics

Before combining the proposed metrics to build a user curiosity model we first perform a careful exploration of each individual metric. In particular, we identify complementary and redundant metrics, before computing our final stimulus. Moreover, we also present a characterization of user profiles based on the subset of metrics that we consider.




Figure 4.2: Fraction of redundant cases for each pair of metrics.

4.3.1 Variable Selection

In order to understand how metrics relate to one another, we measured the Spearman correlation coefficient ρ_s between all possible pairs of metrics. Recall that this coefficient ranges from $[-1, 1]$, with each extreme indicating perfect correlations. The Spearman coefficient was measured by creating a vector for each metric considering every access i of a user. Thus, if a user has accessed 10K songs, we have 7 vectors (one for each metric) of size 10K to correlate.

Our end goal is to filter out redundant metrics. Thus, with the correlation coefficient we aim to identify metrics where for the majority of users the Spearman correlation exceeds a threshold of ± 0.5 . In particular, we determine that two metrics are *redundant* when: over 90% of users have the value ρ_s *outside* of the range: $\rho_s \notin [-0.5, +0.5]$. When correlations are in this range for 90% of the users, we define the pair of metrics as complementary.

A matrix with the fraction of cases for which each pair of metrics was considered redundant in Figure 4.2. The figure shows the correlation considering every event of every user. We define the operation $A \longleftrightarrow B$ as the relationship of correlation between two variables, where A and B are two variables from stimulus metric. It serves as initial evidence of redundant metrics. After inspecting the figure, we computed our redundancy threshold (above). We found evidence redundant relationships as follows:

- For 96% of users: $genreConflict \longleftrightarrow genreUncertainty$;
- For 100% of users: $genreConflict \longleftrightarrow overallGenComplex$;
- For 96% of users: $overallGenComplex \longleftrightarrow genreUncertainty$.

For all other pairs of metrics, the proportion of users for whom they are strongly correlated is less than 35% as shows Figure 4.2, i.e., these pairs are complementaries. This means that these metrics can capture some important information about user's curiosity behavior in online music consumption.

To understand these findings, recall that *overallGenComplex* is related to the number of genres inside window's period of time. *genreConflict* is the probability of genre on the window. In the end, both metrics end up capturing correlated stimuli. Also, *genreUncertainty* is the entropy of genres inside window which, by definition, is derived from *genreConflict*.

Although there are some cases where above correlation are somewhat expected, there are interesting cases of metrics capturing semantically related concepts (e.g., the novelty of artists, songs, and genres) that were not correlated. By filtering out redundant metrics, we can better understand the curiosity of users. Based on these findings, we keep five metrics: *songNovelty*, *artistNovelty*, *genreNovelty*, *instGenComplex* and *genreUncertainty*.

4.3.2 Clustering by Access Profile

We now turn our attention to uncovering the common curiosity patterns in our dataset. According to Section 4.2.3, there are close to 150 million listening events when we consider the 43 thousand users with at least 1,000 listening events. In order to understand the common patterns, we employ a clustering approach. In particular, measure each one of the five *non-redundant* metrics). In the end, we have a 150 million by 5 matrix that we use to cluster *listening events*.

In our clustering, we make use of the Mini-Batch K-Means that that is suitable for large Web datasets [138]. To define the number of clusters k , we made use the β_{CV} and Average Silhouette score. The later, Average Silhouette score, is an aggregate measure of how similar an element is to its own cluster compared another clusters. The β_{CV} is defined as the ratio of coefficient of variation (CV) of intra-clusters distances and CV of inter-cluster distances. The first k value smaller than the smallest value of k at which β_{CV} remains approximately stable should be selected [99].




Figure 4.3: Determining the number of access profile.

Before clustering, we re-normalized the metrics using a Z -scale. It is well known that K-Means will lead to poor clustering when values are not normalized. Figure 4.3a shows the results of the average Silhouette score with a 95% confidence interval (i.e., represented by the error bars and the blue shaded area) as a function of the number of clusters k obtained from 10 runs of clustering experiments. In this figure, the best value of the average Silhouette score is for $k = 2$, where it has the largest value, i.e., almost 0.23 with a narrow interval. However, this value is smaller than 0.25 and therefore requires an additional analysis method [64]. Therefore, we use the β_{CV} method and the results are shown in Figure 4.3b. The smallest value of k at which the β_{CV} starts to remain approximately stable is $k = 5$. The β_{CV} values vary in the interval $[0.8, 1.5]$ from $k = 5$ to $k = 18$, which is a short range of 0.07. Therefore, we choose the first k value smaller than 5, namely $k = 4$, and this is the better choice.

To interpret our clusters, in Figure 4.4 we depict the centroids of each cluster. Recall that a centroid is simply the average of each of the 5 metrics. To provide further evidence that our clustering is adequate, we computed the confidence interval of 95% for each metric from the centroids. Based on these intervals, we verified that each cluster is different from the others (there is no overlap in the intervals for each metric).

To better understand our clusters, we adopt the labels used by [118]. In particular, it is important to understand that each cluster summarizes the type of listening events, not users. These are labelled as follows:

- C_0 (*Underground*): The access is towards an item with similar stimuli metrics. However, the values are low indicating a tendency to stay in the relaxation zone for every metric. There are 48,627,376 events in this cluster and it is exhibited in the Figure 4.4a and represents 32.43% of listening events;
- C_1 (*Eclectic*): Notice that the accesses in this cluster is regularly spread out across each metric. Different from the previous case, the metrics are higher (i.e. higher





Figure 4.4: Radar-chart with normalized stimulus metrics per cluster.

stimulus). 52,829,351 events and it is in the Figure 4.4b with 35.23% of listening events;

- C_2 (*Niche Radical*): the Figure 4.4c exhibits the radar-chart for this cluster, a set of access which the user listening to repeated songs (smaller novelty). Also the songs are more complex. In total, 18,645,990 events and denotes 12.44% of events;
- C_3 (*Explorer*): the Figure 4.4d shows the *Explorer*. The event is very similar to *Eclectic* with 29,846,794 events. However, its main difference is the accesses pursuing more complex music (i.e. songs with many genres) representing 19.90%.

4.3.3 Clustering by User Behavior

To better understand how users switch between different clusters of access profiles at the level of online music consumption, we have created a User Behavior Model Graph (UBMG) [99], which allows us to characterize user behavior. The UBMG is a Markovian model and Figure 4.6 shows four examples of this. In a UBMG, the nodes represent the event of the access profile (i.e., *Underground*, *Eclectic*, *Niche Radical* and *Explorer*) and the edges denote the transitions from one particular access profile to another. The weighting of the edges indicates the probabilities for the occurrence of transition patterns. The sum of all outgoing probabilities for each state is 1. We weight each edge proportionally



(a) Example of UBMG for a user with history of access UEENXXNUUUX.

(b) Determining the number of the UBMG clusters.

Figure 4.5: Example of UBMG and Silhouette index for number of clusters.

to the number of transitions in the clusters for each user.

For example, if we look at the sequence of songs listened to by a user (i.e. the history of access), each consumed song is labeled according to the type of stimulation profile – *Underground*, *Eclectic*, *Niche Radical* and *Explorer* – and we count the number of times the user has made a transition from one profile to another. To illustrate this, we define the symbols for each access profile as follows: (i) U for *Underground*; (ii) E for *Eclectic*; (iii) N for *Niche Radical*; and (iv) X for *Explorer*. So, Figure 4.5a shows an example for the sequence of history of access as follows: UEENXXNUUUX. In this example, if the user is in the *Underground* profile, he has a 50% chance of staying in this state, a 25% chance of switching to the *Eclectic* profile, and a 25% chance of switching to the *Niche Radical* profile. As long as the user is in the *Eclectic* state, they have a 50% chance of staying in this profile and a 50% chance of switching to the *Niche Radical* state.

After creating the UBMGs for each user, we attempt to identify patterns of curiosity stimulation in online music consumption by clustering them. Then, we use Mini-Batch K-Means algorithm and use the Silhouette score to select the best number of clusters. For this purpose, we run clustering experiments with number of clusters $k \in \{2, 3, \dots, 18\}$ and, for each value of k , we execute 10 runs to compute the average Silhouette score with a confidence interval of 95%. Figure 4.5b shows the results of average Silhouette score for each value of k . In this figure we see that the best number k of the cluster is 2, which is the largest value of the Silhouette score of almost 0.30 with a narrow confidence interval. The average values of Silhouette decreases according as the k values increase, also the confidence intervals becomes larger (e.g., after $k = 4$) showing higher variability of the results.

Despite the best number of cluster $k = 2$, we analysis the centroids for the top three best values of k , namely 2, 3 and 4 which have the average values of Silhouette score

Table 4.4: User-level curiosity stimulation profiles.

UBMG state	User-level profiles			
	U_0	U_1	U_2	U_3
<i>Underground</i>	15.97%	37.24%	38.02%	10.13%
<i>Eclectic</i>	67.99%	8.55%	30.11%	32.33%
<i>Niche Radical</i>	1.59%	40.38%	12.34%	8.41%
<i>Explorer</i>	14.44%	13.82%	19.52%	49.12%

of 0.30, 0.28 and 0.26 (vide Figure 4.5b). This top k values were selected to be analysed because they have Silhouette score more than 0.25³ and narrow confidence intervals. After manual inspections of the centroids, we checked that for $k = 4$ the centroids of clusters demonstrate better representation of curiosity stimulation pattern. These patterns are depicted by centroids UBMG U_0 , U_1 , U_2 and U_3 of Figure 4.6, in previous analysis we note that clusters U_0 and U_1 were present in all centroids of $k = 2$ and 3, yet we note that the clusters U_2 is a variant of U_0 and U_3 is a variant of U_1 . Therefore, the number of k chosen was 4 because it represents better the diversity of curiosity stimulation profiles existing for all user's behaviors.

Table 4.4 shows each of the UBMG clusters broken down by access profile, i.e., this table shows the percentages for the Underground, Niche Radical, Eclectic, and Explorer access profiles as a percentage in the rows for each UBMG cluster in the columns. Also, the state transitions between access profiles are shown for each user-level cluster in Figure 4.6. At a first glance, one can clearly see a tendency towards repeated behavior for all clusters in Figure 4.6, i.e., the largest probability transitions belong to self-loops. This is somewhat expected as repeated consumption of songs/artists is common online [15].

In Figure 4.6a, UBMG U_0 represents almost 58% of the curiosity stimulation profiles of users (24,426 users) and also has the largest average number of listenings per user of 4,209. In this case, the proportion of self-loops is the highest in all phases. Nevertheless, the presence of access profiles is the lowest for *Niche Radical* with less than 2% of listening events, *Explorer* and *Underground* account for almost 15% of access profiles, and the most represented access profile is *Eclectic* with almost 68% (see Table 4.4). Even more interestingly, users are more likely to migrate from the *Eclectic* and *Explorer* behaviors, i.e., there is a 14% and 20% trend in migrate probability, respectively. When users are in the *Underground* and *Niche Radical* phases, this tendency to migrate decreases.

UBMG U_1 from Figure 4.6b, unlike UBMG U_0 , has smaller probability transitions for self-loops, except for the *Eclectic* behavior, which has a self-loop of 79%. Of particular interest is the single UBMG that has the largest transition weight, from *Underground* with 28% and from *Explorer* with 58% to *Eclectic*, although *Eclectic* is the least represented access profile with 8.55%. Moreover, the *Niche Radical* state moves to the *Underground*

³According to the literature, this threshold of 0.25 for Silhouette score is an indication of existing good clustering structure [64].

state with a probability of 27%.

We emphasize that there are no transitions from the *Eclectic* and *Explorer* behaviors to *Niche Radical* for UBMG of Figure 4.6b. This means that once users are in these behaviors, they do not return to the behaviors that are more conservative (i.e., repetitive) in terms of music consumption as *Niche Radical*. However, This rarely happens because according to the Table 4.4, *Eclectic* is the least represented state (almost 9%), followed by *Explorer* (almost 14%). Nevertheless, U_1 is the second largest cluster in terms of number of users, 9,046 users, it represents 21% of users and has an average of 2,931 listening events per user.

Figure 4.6c shows users in profile U_2 , which is a variant of U_0 . We can see that self-transitions are strongest in each state. However, for U_2 there are more higher migrations to *Underground* behavior than for U_0 from *Eclectic* and *Explorer* behavior to know 31% and 16%, they are almost twice as high. Another significant difference is that there are 4 times fewer users than U_0 , it contains 6,162 users, which is almost 15%. Music consumption is also slower, i.e. an average of 1,930 listens per user. It is interesting to note that, according to Table 4.4, U_2 has a balanced presence of access profiles with 50% of listening events of more conservative type (e.g., *Underground* and *Niche Radical*), that is, more repetitive music consumption, while it has the presence of less repetitive access profiles with almost 50%, i.e., *Eclectic* and *Explorer*.

Finally, Figure 4.6d shows cluster U_3 where users have the second highest music consumption, i.e., they consume on average 3,048 songs per user, although it is the smallest in terms of number of users, accounting for only 7% of users (2,878 users). There are large migrations from all behaviors to *Explorer* and this is a unique feature compared to other clusters. For example: (a) from the state *Underground* there is a probability of 34% to migrate to the *Explorer* state; (b) a probability of 44% to migrate from the *Eclectic* state to the *Explorer* state; finally (c) a probability of 29% to migrate from the *Niche Radical* state to the *Explorer* state.

We recall that *Explorer* has a tendency to consume songs with many genres and is thus distinct from other access profiles. Thus, in Figure Figure 4.6d, users of this cluster tend to consuming more complex musical genres and behave less repetitively when consuming music. In fact, the access profiles for *Eclectic* and *Explorer* add up to 81% and dominate user behavior (see Table 4.4). We also recall that U_3 is a variant of U_1 and thus there are large migrations between the *Eclectic* behavior and the *Explorer* behavior, i.e., more than 26%.

So far, our results point out that user behavior in terms of curiosity is quite complex. Each access profile represents a different pattern of behavior (metric values). Also, users tend to transition across these patterns showing different curiosity stimulation profiles at user level. These results motivate the need of Wundt models that are able to capture such complexities. These models are discussed in the next section.




Figure 4.6: User-level curiosity stimulation profiles: state transition diagrams.

4.4 Curiosity Models

We now tackle our second research question: To which extent user curiosity can be accurately modelled by a Wundt's curve? Following the approach of [186], we can estimate a *single* Wundt's curve as a Beta distribution. The Beta is a continuous probability distribution with shape parameters α and β . The probability density function (PDF) of Beta distribution is:

$$b_{\alpha,\beta}(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}, \quad \text{where } B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$$

and Γ is the Gamma function, α and β are two positive shape parameters. The Beta is a flexible distribution that is able to capture different shapes related to Wundt's curve which is an inverse U -shape and can be estimated with Beta parameters $\alpha > 1$ and $\beta > 1$. In

Table 4.5: Validation of users.

Description	Users (Percentage)
1 Beta	24,578 (57.82%)
2 Beta mixture	15,498 (36.46%)
3 Beta mixture	1,684 (3.96%)
No valid users	751 (1.77%)
Total of users	42,511 (100.00%)

our analysis, the accesses of each user will be used to learn multiple Wundt curves. That is, we shall estimate Beta distributions for each user using as input the access metrics.

To fit the Beta models to our data, we need to map the final value of curiosity to the $[0, 1]$ range (the domain of the Beta). Thus, we initially take the mean of the 5 metrics (notice that they are all in bits) for each listening event. Thus, we are capturing the average number of bits of the access. We then map the mean to $[0, 1]$ range using a min-max normalization.

One of the fallacies of [186] was to assume that each user may be modelled by a single Wundt's curve. Here, we shall show that this is not the case. To capture more complex behavior, we employ a mixture of Beta distributions, i.e.:

$$f(x) = \pi_1 b_{\alpha_1, \beta_1}(x) + \dots + \pi_M b_{\alpha_M, \beta_M}(x) = \sum_{i=1}^M \pi_i b_{\alpha_i, \beta_i}(x), \quad (4.9)$$

where M is the number of mixture components and π_i, α_i and β_i are respectively the i -th mixture component and two shape parameters from Beta mixture distribution. We estimate mixture of betas considering $M \in \{1, 2, 3\}$. For the estimation of parameters, we used the Iterated Method of Moments [137].

4.4.1 Classes of Wundt's Curve

For each user in our dataset, we fit Beta mixtures considering $M \in \{1, 2, 3\}$. To validate each mixture, we employed the Kolmogorov-Smirnov (KS) Test. Notice that we are fitting thousands of users, and thus are performing thousands of hypothesis (one for each user) tests. Given that this is a multiple comparison setting, to correct our p -values, we employ the Benjamini-Hochberg (BH) method. After p -value correction, we employ a *corrected* significance level of $\alpha^+ = 0.01$. We state that each user is modelled by the smallest mixture of Betas, smallest M , which is validated by the test.

So, when we consider all of the users, in Table 4.5 we show the number and respective percentage of users validated for $M \in \{1, 2, 3\}$ mixture of Betas. From table 57.85%

are accurately modelled by a single Beta, 36.46% of users require a two Beta mixture, and 3.96% of users are modelled by three Betas. In the end, only 1.76% were not validated by at least three Betas. These users were discarded. To exemplify our results, in Figure 4.7 we show four examples of users that were modelled with two and three mixtures. The Figures 4.7a, 4.7b and 4.7d clearly show the multi-modality of curves.




Figure 4.7: Mixed Wundt's curve: examples of 4 users.

4.4.2 Curiosity and Access Patterns

We also correlated the mixture of Betas models with the types of access (clusters). To do so, Table 4.6 exhibits the fraction per access profile inside of each mixture model. This fraction is computed over all of the accesses for every user captured with 1, 2 or 3 mixtures (each column sums to 100%). From the table, we can see that *Underground*

Table 4.6: Statistics of Access Profile values per Beta (mixture) distribution.

Access Profile	1 Beta	2 Beta mix	3 Beta mix
<i>Underground</i>	35.54%	35.09%	36.11%
<i>Eclectic</i>	34.39%	32.39%	20.49%
<i>Niche Radical</i>	10.70%	15.38%	27.47%
<i>Explorer</i>	20.37%	17.14%	15.93%
Total:	100.00%	100.00%	100.00%




Figure 4.8: PDF of Access Profile per Beta distribution.

has almost constant presence in each different model (rows). In contrast, *Eclectic* and *Explorer* show a decrease in presence as the number of mixture increases. This indicates that users captured with a single beta will mostly be stimulated by *Underground* and *Eclectic* accesses. More complex users, captured more (2 or 3) Betas, show a higher variation in such behavior. Overall, these findings serve as evidence that users that transition across multiple clusters will require more mixtures.

To exemplify these findings, Figure 4.8 we show the density of accesses in each mixture model. In Figure 4.8a we show users captured by a single Beta, Figure 4.8b represents those captured by a two Beta mixture, and Figure 4.8c a three Beta mixture. From the figures, notice that the frequency of *Niche Radical* access profile is clearly increasing when the number of mixtures increases. The users in Figure 4.8a are more stable, being curious towards *Underground* and *Eclectic* stimuli. A increase in value will lead to a new Beta distribution. Specially when it happens around a different stimulus (x -axis).

Complementing Figure 4.8, Figure 4.9 shows the four examples from Figure 4.7 broken down per access profile. Both the overall figure and the one per user show similar patterns. That is, in cases 4.9a, 4.9c and 4.9d is easy to see that accesses are concentrated in with *Underground* and *Niche Radical* accesses. On the Figure 4.9b the main concentration areas involves *Underground*, *Eclectic* and *Explorer*, one can to see that greater mode (i.e. more than 0.5) is skewed to the right of Figure 4.7b.

Finally, to better understand the relationship between the mixture model and clusters when considering each user, in Figure 4.10 we show the Cumulative Distribution Function (CDF) of the proportion of accesses *per cluster* when *normalized per user*. That




Figure 4.9: Access Profile: examples of four users.

is, to compute the CDF, each data point was the fraction of accesses in a cluster when normalized considering a single user only (e.g., each point captures a user and each curve a cluster). Figure 4.10a depicts the CDF for the users modeled with a single Beta, whereas Figures 4.10b and 4.10c shows the fractions for the users captured with M equal to 2 and 3.

Similar to the overall behavior, notice how the *Niche Radical* is less present when considering users modeled with a single Beta. Also, we can see that the *Eclectic* curve becomes less present. In particular, consider that the median of each curve (y -axis equal to 0.5) is of 0.34 for *Underground*, 0.31 for *Eclectic*, 0.05 for *Niche Radical*, and 0.17 for *Explorer*, when we have a single Beta. That is, for 50% of the users modeled with a single Beta we have over 31% of accesses in the *Eclectic* cluster. For *Niche Radical* the median is only 5% of accesses. Contrast this setting to the one when we have three Betas. The medians are of: 0.36 for *Underground*, 0.15 for *Eclectic*, 0.22 for *Niche Radical*, and 0.12




Figure 4.10: CDF of fraction of accesses per profile per user.




Figure 4.11: Division of the user's activity period into short time period (e.g., weeks).

for *Explorer*. In the end, there is a significant increase for *Niche Radical*. There is also a large decrease in the *Eclectic* accesses. This result shows not only the events overall are complex, but also that user behavior is complex.

Overall, these findings combined with Figure 4.8 serve as evidence that users that transition across multiple clusters will require more mixtures. In other words, users with more complex behavior require more complex models, leading to decline the fallacy of a single Wundt curve.

4.4.3 Temporal Analysis

In this section, we examine the temporal analysis of Wundt's curve of users estimated with 2 Beta and 3 Beta mixture. Thus, we investigate how the user's curiosity profile modeled by a mixture of Betas evolves over time, i.e., whether the user's curiosity profile is always estimated by a mixture of Betas even for short time periods. For this purpose, we select the 25 users with the most listening events for each of the two types of Beta mixtures (50 users in total), i.e., 2-Beta and 3-Beta Mixture for Wundt's curve, respectively. The 50 selected users belong to the 99th percentile of the number of accesses (i.e., listening events) per user. The 25 users for the 2-Beta mixture therefore account for 0.16% of 15,498 users and the 25 users for the 3-Beta mixture account for 1.5% of 1,684 users (see Table 4.5). In addition, the selected users have between 33,000 and 174,000 listening events. This is important as these users have access throughout the analysis period (i.e., approximately 1 year).

The analysis consists of dividing the user's every activity period (at most a little more than 1 year) into short time periods (e.g., days, weeks or months) and estimating Wundt's curve (the user's curiosity profile) for each of these time periods for those with

at least 30 listening events. Figure 4.11 shows these procedure where: Figure 4.11 (a) represents an example of an user's activity period in the dataset; Figure 4.11 (b) shows how we divide the entire period of activity into short time periods, in this case one week; and Figure 4.11 (c) shows that for each time period we estimate Wundt's curve to check whether it is estimated by a mixture of Beta distributions (2 or 3 components) or by a single Beta distribution.

To evaluate our experiments, we define two metrics, namely $avgPrecision_\beta$ and $avgPrecision_{AP}$, which represent *the average precision of estimate Beta Mixture distribution* and *the average precision of access profile*, respectively. These metrics are defined as follows. First, we need to start with the metric *precision of estimated Beta Mixture distributions*, which is defined as the proportion of total number of user's short time period which Wundt's curve were estimated by a mixture of Beta distributions (2 or 3 components) for a user. It is defined as

$$Precision_\beta = \frac{TB}{TP}, \quad (4.10)$$

where TB is the total number of short time periods with Wundt's curve estimated by a Beta mixture distribution and TP is the total number of short time periods. The *average precision of the estimated Beta Mixture distribution* is computed as the average over all users for estimated Beta Mixture distributions and it is denoted by $avgPrecision_\beta$.

The metric *precision of the access profile*, in turn, is defined as the proportion of the total number of short time periods that have more than one access profile for a user. We have therefore defined it as

$$Precision_{AP} = \frac{TAP}{TP}, \quad (4.11)$$

where TAP is the total number of short time periods with different access profiles (i.e., at least two types) and TP is the total number of short time periods. Then, *the average precision of the access profile* per user is calculated as an average over all users and denoted by $avgPrecision_{AP}$.

Each of the metrics is in the range $[0, 1]$. If the value of the metric is close to 1, this is an indication that the majority of users their short time periods with Beta mixture distribution or that they have different types of access profiles in their short time periods, i.e., they are very dynamic in terms of their access profile.

Table 4.7: Results of experiments for temporal analysis with I.C. of 95%.

Metrics	2 Beta mixture	3 Beta mixture
$avgPrecision_\beta$	0.972 ± 0.023	0.920 ± 0.043
$avgPrecision_{AP}$	1.0	0.994 ± 0.010

Table 4.7 shows the results of experiments according to each of the metrics mentioned above. In this table, it can be seen that for all metrics, more than 90% of users have




Figure 4.12: Example of user with 2-Beta mixture: Wundt's curve and access profile.

a more complex curiosity profile with more than a single Beta distribution for Wundt's curve even for short time periods. Nevertheless, almost 100% of the time periods contain a different stimulus profile of curiosity, even for short time periods. These results show that users are very dynamic when consuming online music.

Figure 4.12 and Figure 4.13 show two examples of users with a set of short time period of one week, one user with a Wundt's curve estimated with 2 Beta mixture distribution and one user with a 3 Beta mixture distribution. The figures are arranged in such a way that for each Wundt's curve there is a histogram of the respective access profile underneath. In addition, 12 weeks are shown in these figures for reasons of space.

The user in Figure 4.12 has 49 periods of a week and $Precision_{\beta} = 0.98$ (98%), which means that there are 48 periods out of 49 estimated with a Beta mixture distribution. In addition, this user has $Precision_{AP} = 1.00$ (100%) and it shows that this user is very dynamic in terms of access profiles in all periods of a week.

Finally, Figure 4.13 shows the curiosity profile of a user estimated with 3 Beta mixture distributions. In total, he has 40 periods of one week, in 39 of them Wundt's curve was estimated with a mixture of Betas, i.e., $Precision_{\beta} = 0.98$, in turn, he also has 39 periods with at least two types of access profiles, i.e., $Precision_{AP} = 0.98$, which means that this user has a stable access profile in only one week.




Figure 4.13: Example of user with 3-Beta mixture: Wundt's curve and access profile.

4.5 Validation of Curiosity Model

Here we use null models to validate our proposed metrics for the collative variables from Section 4.3. To this end, we state two hypotheses that are tested using the null models. These hypotheses are based on the three fundamental assumptions defined in Chapter 3, namely (1) that curiosity can be triggered differently depending on the target platform (i.e., LastFM) and user actions (i.e., music consumption); (2) that curiosity can change over time; and (3) that curiosity has a period of activation – the window of interaction. The hypotheses are defined as described below:

Null hypothesis 4.1 (H4.1): *Curiosity is influenced by user's personal interest and by temporal locality of consumed items – user's tendency to access items repeatedly over time.*

This hypothesis assumes that there is a relationship between the stimulus values and the user's association of time and content item and evaluates this relationship.

Null hypothesis 4.2 (H4.2): *Curiosity is influenced by the frequency and total number of accessed items within the window of interaction.*




Figure 4.14: Design of Null Model 0 and 1.

With this hypothesis, we can assess whether the stimulus values are caused by chance or not.

To avoid long execution times for the validation experiments with the null models, we selected the first 100 users whose total number of accesses is in the 99th percentile of the number of listening events per user, i.e. users with more than 33 thousand listening events. From these users, we randomly select uniformly 100 accesses of each user to compute the stimulus according to our metrics for the null models, which we describe below.

The design of the *Null Model 0* and *Null Model 1* can be seen in Figure 4.14. The figure shows the current song listened by a particular user in the blue shaded area on the right side. This current song is preserved in the null model, while that the songs for the window of interaction are randomly selected from the set of songs listened by the user. This random selection can be done in two ways: (i) for the *Null Model 0*, the songs for the window of interaction are selected according to the distribution of the user's songs preferences, for this purpose, this distribution is computed based on the frequency of all the songs listened by the user; (ii) in the *Null Model 1*, the songs are selected for the window of interaction according to the uniform distribution.

Although the songs are randomly selected, the timestamps are the same for both null models. This is illustrated in the Figure 4.14 by the dashed vertical lines connecting the timestamps of the real sample, represented by the horizontal black line at the top, and the random sample of the null model, represented by the horizontal red line at the bottom. In summary, *Null Model 0* aims to break the *association* between the window of interaction and the listened songs by randomly selecting songs for the window of interaction, while *Null Model 1* aims to break the *association* between the listened songs and the *user's personal interest* by selecting songs according to a uniform distribution, which is far from the user's interests. In this way, these null models were developed to test **H4.1**.

With the goal of testing **H4.2**, the *Null Model 2* was developed following design




Figure 4.15: Design of Null Model 2.

of the Figure 4.15 as an extension of *Null Model 1*. In this figure, we can see that the timestamps are not retained and it is the difference with the other model, in contrast, they are randomly selected from the time interval of the window of interaction. Also, the number of songs in the window of interaction is different. This is the main difference of these null model, the number of songs is randomly generated in accordance with the user's preference. So, this second model assesses that stimulus values is not caused by chance.

The distribution of the number of songs is estimated according to the frequency of the total number of songs in each window of interaction of user, whereupon, a number of songs is randomly selected from this distribution. From this number of songs, the songs that compose the window of interaction are uniformly selected according to the set of songs of the user – which is extracted from his historical of listening events. Next, the timestamps for each selected song should be randomly chosen within the time interval of the window of interaction. So they are generated evenly after the previously selected songs, i.e., one timestamp for each song. The Figure 4.15 shows that in the window of interaction of the real sample at the top, there are 9 songs and 4 different types of songs (represented by the play icons in the colors red, blue, purple and yellow). While in random sample at the bottom of figure, there are 12 songs and 6 different types of songs (represented by the play icons in the colors orange, blue, red, purple, black and green). Also, we can see that the timestamp for the current song is preserved for both samples in this figure. This means that the timestamp t_{10} of the real sample above and the timestamp t_{13} of the random sample below are identical.

In short, the *Null Model 2* breaks the *frequency of songs* and the *total number of songs* that have an impact on the user's curiosity, i.e., it tests **H4.2**. Since this null model is an extension of the *Null Model 1*, the uniform selection of songs also breaks the *personal interest* and the *temporal constraints* of user (after randomly selecting each timestamp of songs), which represents the test of **H4.1**.

To test each of the previously defined hypotheses, the difference between the real




Figure 4.16: CDFs of the differences per access profile for *artistNovelty*.

stimulus values and the random stimulus values was calculated. In this way, positive values for this difference indicate that the real stimulus values are greater than the random stimulus values for the null model. Negative values, on the other hand, indicate that the random stimulus values are greater than the real stimulus values. So, the further the null models are from the real behavior of the users (i.e., from the real stimulus values), the larger the difference between the random sample and the real sample. This means that surprisal and entropy of the metrics are larger than the measured values of the real users.

Next, we plot the cumulative distribution functions (CDFs) of the differences between the real stimulus values and the random stimulus values from the null models for the metrics⁴ *songNovelty*, *artistNovelty*, *genreNovelty* and *genreUncertainty* per access profile (i.e., *Underground*, *Eclectic*, *Niche Radical*, and *Explorer* presented in Section 4.3.2) of 100 selected users and their respective 100 accesses. The Figures 4.16 show the CDFs for the *artistNovelty* metric for the *Null Model 0*, the *Null Model 1* and the *Null Model 2*. Similarly, the CDFs of the differences for the *songNovelty* metric are shown in Figures 4.17. The Figures 4.18 again show the CDFs of the differences for the *genreNovelty* metric. Finally, the CDFs of the differences for the *genreUncertainty* metric are shown in Figures 4.19.

Regardless of the null model, we can see that the distributions in the three plots in Figure 4.16, Figure 4.17, Figure 4.18 and Figure 4.19 are very different. The values of the difference distributions are also mostly negative, i.e. the random samples give the largest metric values, and from the largest to the smallest distribution we have the following access profiles: *Niche Radical*, *Underground*, *Explorer* and *Eclectic*. For the *genreNovelty* in Figure 4.18, the only difference between these sequences is that *Eclectic* is larger than *Explorer*. Although for all metrics, the access profiles with repetitive behavior, such as *Niche Radical* and *Underground*, have the largest difference distributions and the access profiles with less repetitive consumption behavior, i.e. unexpected patterns (*Eclectic* and *Explorer*), have the smallest difference distributions.

⁴It is important to note that the null models do not change the current song according to the above description of the null models, because only the metric *instComplexity* has no CDFs of differences, as




Figure 4.17: CDFs of the differences per access profile for *songNovelty*.

For example, if we look at the *artistNovelty* metric in Figure 4.16 for the *Niche Radical* access profile (red dashed line), the median for *Null Model 0* of Figure 4.16a is almost -4 and -5.75 for *Null Model 1* and *2* in Figures 4.16b and 4.16c respectively. Similarly, the median of *Underground* (blue solid line) is -3.5 for *Null Model 0* in 4.16a, -4 for *Null Model 1* in Figure 4.16b and *Null Model 2* of Figure 4.16c. Now, looking at the *Explorer* access profile for the metric *artistNovelty*, the median of differences is -1 for *Null Model 0* in Figure 4.16a, -1.8 for *Null Model 1* in Figure 4.16b and -2 for *Null Model 2* in Figure 4.16c. Similarly, in the same figures, for the *Eclectic* access profile, the median is 0 for *Null Model 0* and *Null Model 1*, at last, the median is -1 for *Null Model 2*.

For *songNovelty* in Figure 4.17, the median of *Niche Radical* is -3 for *Null Model 0* in Figure 4.17a, -3.8 for *Null Model 1* in Figure 4.17b and -3.9 for *Null Model 2* in Figure 4.17b. Furthermore, for the *Underground* access profile, the median is 0 for *Null Model 0* and *1* in Figures 4.17a and 4.17b respectively. For *Null Model 2* in Figure 4.17c the median is -1 , moreover, this is the only plot for the *songNovelty* metric with an S-shape curve and shows that *Null Model 2* has the largest difference. It can be seen that for the other null models, more than half of differences are zero. If we look at the *Explorer* access profile, we can see that for *Null Model 0*, *Null Model 1* and *Null Model 2*, up to 17%, 24% and 25% of the differences between users have values less than or equal to -1 . Already for the *Eclectic* access profile, up to 6%, 8% and 50% of the differences between users have values less than or equal to 0 for *Null Model 0*, *Null Model 1* and *Null Model 2* respectively, showing the tendency for differences to increase when switching from one null model to another.

Figure 4.18 shows the CDFs for the *genreNovelty*. For the *Niche Radical* in the *Null Model 0* in Figure 4.18a, up to 30% of the differences of users have a value of less than or equal to -1.2 . For *Null Model 1* and *Null Model 2*, however, the percentage has a slight slope for the same value of the difference and increases up to 38% of differences. On the other hand, the differences between the distributions of the null models are smaller its values remain the same as the current songs.




Figure 4.18: CDFs of the differences per access profile for *genreNovelty*.

for the *Underground* access profile than for *Niche Radical*. For example, up to 75% of differences the values are less than or equal to -0.07 for the *Null Model 0* in Figure 4.18a, while the value of differences for *Null Model 1* and *2* is -0.13 for the same percentage. In the case of the *Explorer* access profile, as can be seen in Figure 4.18a for the *Null Model 0*, the differences are negative for up to 56% of differences, while the curves for the *Null Model 1* in Figure 4.18b and the *Null Model 2* in Figure 4.18c are identical and up to 50% of differences have negative values. In the *Eclectic* access profile, unlike other CDFs of this profile for the same metric, all curves for the null models are identical. We recall that for this profile the entropy of musical genres is the highest, which means that its behavior is closest to the random pattern generated by all null models.

Last but not least, for the access profiles with repetitive behavior, the CDFs of the differences for the *genreUncertainty* metric are shown in Figure 4.19. For the *Niche Radical* access profile and *Null Model 0*, the median is -1 , while the median for *Null Model 1* and *Null Model 2* is -1.4 . Finally, the median of *Null Model 0* is -0.4 and the median of *Null Model 1* and *Null Model 2* is -0.53 for the *Underground* access profile. The profile *Explorer* shows a median of -0.2 for the *Null Model 0*, -0.4 for the *Null Model 1* and -0.5 for the *Null Model 2*. And for the profile *Eclectic*, a median of -0.1 , -0.2 and -0.3 can be seen for the *Null Model 0*, the *Null Model 1* and the *Null Model 2* respectively.

In summary, for access profiles with repetitive consumption behavior (i.e., predictable behavior), such as *Niche Radical* and *Underground*, frequency of access dominates the window of interaction and thus surprise or entropy is greater and the difference is larger. In contrast, for access profiles with less repetitive consumption behavior (i.e., less predictable behavior), such as *Eclectic* and *Explorer*, the total number of accesses dominates the window of interaction and thus surprise or entropy is smaller and the difference is lower.

As a result, our three proposed null models test **H4.1** and **H4.2** and show that by breaking these hypotheses, we can confirm that the stimulus values from null models are larger than the real stimulus values, which is an indication that the measured stimulus




Figure 4.19: CDFs of the differences per access profile for *genreUncertainty*.

values of the users in the LastFM are not caused by chance.

4.6 Summary

In this work we provide an in-depth analysis on how to model curiosity in online content consumption. Motivated by advances both in data mining as well as psychology, we model curiosity as a mixture of Beta distribution. Each distribution capturing a Wundt’s curve. In particular, we answer two research questions:

- **RQ1:** *Are there distinct user behavior profiles in terms of curiosity stimuli, as captured by the collative variables?*
- **RQ3:** *To which extent user curiosity can be accurately modeled by a Wundt’s curve?*

Based on the methodology of Berlyne [16], who proposes the use of the concepts of surprisal and entropy from the field of information theory [92, 32] to quantify the curiosity of individuals, we have proposed metrics to investigate the four aspects of collative variables such as novelty, uncertainty, conflict and complexity. Our results are important because they reveal some of the weaknesses of previous efforts. In particular, we carefully examine which collative variables should be used to model curiosity stimuli. We therefore endeavor to adapt the key concepts of Berlyne’s methodology for LastFM, an online music consumption platform. To this end, based on these fundamental concepts, we proposed seven metrics to capture these four aspects of collative variables in the context of LastFM, namely: (1) novelty of song; (2) novelty of artist; (3) novelty of genre; (4) uncertainty of music genre; (5) conflict of music genre; (6) instantaneous complexity of music genre; and (7) overall complexity of music genre.

In this context, we show that user curiosity is more complex than it seems, as the same user can be curious about different profiles (clusters) of stimuli (RQ1). To find

out how these metrics relate to each other, we evaluate the correlation between them for each user. Thus, we find a reduced set of metrics that complement each other (i.e., song novelty, artist novelty, genre novelty, instantaneous complexity of music genre). With this reduced set of metrics, we can discover four different access profiles for stimulating curiosity at the level of song consumption. To characterise these access profiles, we used the specific nomenclatures for online music consumption of [118] and named them as follows: *Underground*; *Eclectic*; *Niche Radical*; and *Explorer*.

To better understand how users switch between different access profiles, we characterized user behavior clustering the User Behavior Model Graph. We also found that user behavior is very dynamic and is composed of different access profiles. Therefore, we identified four different user behaviors: (i) the UBMG U_0 comprises 57.5% of users and when these users are in a certain access profile, they tend to stay in it; (ii) the UBMG U_1 , which comprises 21.28% of users who tend to migrate from the *Niche Radical* and *Explorer* access profiles to *Underground* and *Eclectic* respectively; (iii) the UBMG U_2 has 14.5% of users and is a variant of U_0 , but tends to migrate from the *Eclectic* and *Explorer* access profiles to *Underground*; finally (iv) the UBMG U_3 has the lowest percentage of users, 6.77%, moreover, its users have great chances to switch between *Eclectic* and *Explorer* access profiles, users also tends to migrate from *Niche Radical* to *Explorer*.

Towards to answer RQ3, we find that for a large proportion of users, more than one Wundt curve is required to model curiosity. This means that we need 2 or 3 beta mixture distributions to model the curiosity profile for more than 40% of users. In addition, it is worth noting that these users have different levels of access profiles. For example, the distribution of *Eclectic* per user decreasing from 2 Beta mixture models to 3 Beta mixture models. Recall that this profile is characterised by high stimulus values in terms of novelty of music genre, artist, and song, and high diversity in terms of genres (i.e., uncertainty and complexity). In contrast, the profile *Niche Radical* shows an increase in its distributions per user from 2 Beta mixture models to 3 Beta mixture models and it shows an increase in the profile that has inherently low stimulus values for its access profile. These results emphasize the complexity of curiosity profiles and it also shows the obvious shortcomings of previous studies that did not provide an accurate fit of Wundt’s curve because they used a unique beta distribution to model the user’s curiosity profile.

The complexity of the curiosity profile becomes even clearer when we examine the temporal analysis of Wundt’s curve. We then examine how the curiosity profile evolves over time when we estimate Wundt curve by mixture of Beta distributions. To this end, we conducted experiments in which users’ activity periods were divided into short time periods (i.e., from years to weeks). In this way, we can check that for each piece of the short period is still estimated by a mixture of Beta distributions, i.e., 2 or 3 Beta mixtures. Not surprisingly, in more than 92% of the cases, the user’s curiosity profile for each short time period can be estimated by a mixture of Beta distributions. Moreover, in more than

99% of cases, curiosity profiles are composed of different access profiles, i.e., more than two types of curiosity stimulation during song consumption.

Finally, to validate our metrics, we use null models based on the three fundamental assumptions for the construction of our metrics (defined in the Section 3.2), so, they are random sample generations derived from the user's history of listening events. Thus, we explore the fact that the further away the null models are from the users' real behavior, the greater the difference between the users' real measurements. Interestingly, the results show that for access profiles with predictable behavior, such as *Niche Radical* and *Underground*, the frequency of access dominates the sample of null models and has a greater impact on the values of the metrics. Therefore, the difference between the metrics of the null models and the real metrics of the users is larger. In contrast, for the profiles *Eclectic* and *Explorer*, whose behavior is less predictable and dominated by the total number of access, the metrics show a smaller difference to the real measures of users. In summary, these results show that our metrics can indeed capture the stimuli of curiosity of user behavior in online music consumption.

As a limitation of this results, we emphasize that our analysis covers only what the data can provide. For example, the musical genre of the song is also assumed to be the genre of the artist. In fact, an artist can be classified by some musical genres, which can be different from the musical genres of his songs. In addition, the artist may record a song with other artists who have different musical genres (i.e. styles), so the musical genres of the song may also differ from the artist's style. This limitation is due to the fact that the musical genres available from [132] for the dataset of [131] are limited and we are therefore restricted. Moreover, specific aspects of the platform and exogenous effects of the real world are not included. The former refers to song recommendations, ratings of songs, artists or albums, reviews, social sharing (e.g. playlists and songs, collaborative playlists and group sessions), redirects from other platforms (online social networks, websites, communication and messages, and streaming services) and finally applications developed via the platform's Web API. The exogenous factors in turn refer to the influence of releases such as albums, singles, videos, remixes and live performances, shows and events as well as the popularity of songs and artists.

Thus, this chapter serves as a basis for designing user-centric applications that consider curiosity. For instance, by identifying whether the user is in anxiety or boredom zone it is possible to stimulate the user to move until the curiosity zone recommending the right items. So, the user will achieve higher levels of curiosity score, hence leading to the user to involve more deeply into system. As future work, we can extend the findings in this chapter to provide better personalized search results and recommendations. Finally, our results can be naturally extended to other domains besides music since that there is some way to classify the set of items on categories. Thus, it will be possible to compute the stimulus metrics according to collative variables and to fit the user's curiosity.

Chapter 5

Metrics of Social Curiosity: The WhatsApp Case

This chapter is organized as follows. Section 5.2 presents initial considerations about the work of this chapter. Next, Section 5.3 describes the WhatsApp communication platform. Our novel metrics of curiosity are introduced in Section 5.4. Our characterization results, including the dataset used, are discussed in Section 5.5. Curiosity models to represent human behavior in a content sharing environment are discussed in the Section 5.6. In Section 5.7 we describe our approach to validate the metrics for social curiosity. We discuss some limitations of our study and their implications in Section 5.8. Finally, summary and future work are offered in Section 5.9.

5.1 Introduction

As discussed in Chapter 2, social curiosity entails two distinct aspects: a general interest in gaining new information about how others think, behave, or act, and an interest in interpersonal information gained through exploratory behavior. Therefore, we believe that social influence can be considered as another collative variable that stimulates curiosity and ultimately determines individuals' behavior.

In this context, our present effort is driven by three research questions:

- **Research Question 2.1 (RQ2.1):** *How to quantify social influence as a stimulus to one's curiosity driving the information dissemination process?*
- **Research Question 2.2 (RQ2.2):** *How does social influence relate to other collative variables priorly associated with curiosity stimulation?*
- **Research Question 2.3 (RQ2.3):** *How are users characterized in terms of social stimulation to curiosity?*

We tackle these questions using WhatsApp as case study. WhatsApp is a free messaging app that has surpassed the mark of 2 billion users worldwide in 2020¹. It connects users in end-to-end as well as group conversations, and the latter has been shown to be an effective vehicle for information dissemination at large [122]. Thus, we here propose to *model the social stimulus to the curiosity that drives users to share content (thus communicating with each other) in WhatsApp groups*.

To that end, we start by proposing a set of novel metrics to quantify such stimulus (RQ2.1). Towards deriving these metrics, we follow the seminal work by Berlyne [16], which proposes a methodology to quantify collative variables related to curiosity stimulation based on information theoretical metrics. Berlyne applied these concepts to the derivation of metrics related to more traditional collative variables (e.g., novelty, conflict). We here follow his arguments, as well as similar ones by other authors [144], and apply his methodology to derive metrics capturing the effects of social influence as a component of curiosity stimulation.

Moreover, following the same methodology and building on our prior experience [147], we also adapt metrics capturing other collative variables related to curiosity stimulation, notably novelty, complexity, conflict and uncertainty to the particular domain of WhatsApp groups. These variables have been studied before in other domains [186, 178, 171, 147] but not in a platform of group communication, such as the one provided by WhatsApp.

Our study is carried out in a dataset consisting of over 2 million messages shared by more than 7.5 thousand users in 335 publicly accessible groups in Brazil during a one-month period. By correlating the two sets of metrics – metrics related to social curiosity and metrics related to the other collative variables – we show that our social curiosity metrics do indeed capture novel aspects of one’s curiosity stimulation which are not represented by the other (priorly studied) collative variables (RQ2.2).

Finally, we use our metrics to characterize the social curiosity stimulation of users when sharing content in WhatsApp groups (RQ2.3). Our characterization is performed at three levels of aggregations: we quantify social curiosity stimulation driving the sharing of individual messages, the overall behavior of individual users as well as WhatsApp groups.

Our main contributions can be summarized as follows:

- We propose four new metrics to capture social influence as a component of curiosity stimulation at the user level, as well as one metric to capture the same aspect at the group level. We instantiate the proposed metrics for the specific context of content sharing in WhatsApp groups, but they are general enough to be applied to other platforms of online group communication. Additionally, we also propose seven metrics that capture other traditionally studied collative variables, namely novelty,

¹<https://www.whatsapp.com/about/>

complexity, conflict and uncertainty, to the same context. To our knowledge, this is the first effort to propose metrics to quantify components of curiosity stimulation in such environment.

- We offer an extensive evaluation of social influence as a component of the curiosity stimulation driving content sharing in WhatsApp groups. Our study reveals that social influence, as captured by our proposed metrics, is complementary to the other traditionally analyzed collative variables, thus reflecting a novel and important component of the curiosity stimulation process. Also, we found great diversity and dynamics in social curiosity stimulation in WhatsApp, at the message, individual and group levels. Specifically, we found three profiles of social curiosity stimulation driving user behavior when sharing individual messages, and used them to uncover five different profiles describing how the social curiosity of a user in a given group evolves over time. We also found evidence that the social curiosity stimulation of a particular user may be quite different depending on the group she participates in, which hints at a role the group has on its members. Yet, different groups may exhibit quite different and very dynamic group-level social curiosity stimulation, as such dynamics results from the aggregation of the behavior of those users that are more actively sharing content at each time.

5.2 Initial Considerations

In this chapter, in addition to the four aforementioned collative variables in the Section 2.1, social influence should also be considered as part of the stimulus to one's curiosity, especially in online social networks and social media applications, where, to a large extent, one's behavior is driven by connections, common interests with others and social observation. This argument is aligned with the concept of *social curiosity*, recently introduced in [63, 61], as a key component of a five dimensional curiosity model. However, these studies were based on surveys with real people.

Here, we are interested in proposing quantifiable measures of social curiosity as a driving force behind online behavior. Specifically, we propose different metrics to quantify the impact of social influence as a component of curiosity stimulation driving users to share content on WhatsApp. To our knowledge, prior attempts to operationalize the concept of social curiosity by proposing quantitative metrics to estimate it are quite scarce and mostly limited [186, 178, 147]. This is probably due to the challenges to capture, quantitatively, the aspects associated with such a highly subjective concept, as argued by

several recent studies in the Psychology and Neuroscience domains [63, 61, 160, 4, 72, 7]. The derivation of our metrics, discussed in Section 5.4, follows the rationale on Berlyne’s original arguments and his proposed methodology, here adapted to capture social curiosity. Specifically, we explore metrics from information theory to estimate how such highly subjective concept can compose the curiosity stimulation process in the particular setup of interest, i.e., group communication on WhatsApp. To our knowledge, we are the first to operationalize Berlyne’s ideas to estimate social curiosity, especially in the selected domain.

The main research gap we address in this chapter with respect to the studies in Section 2.2.2 is the modeling of a novel collative variable associated with curiosity stimulation, namely social curiosity. Thus, the closest study to ours is that by Xu *et al.* [178]. However, unlike this prior work, we do not aim at improving item recommendation. Rather, our primary focus is on modeling user behavior, with particular interest in how users’ decisions to share content is driven by social curiosity. Moreover, we explore a completely different setting – content sharing in a group communication application, where curiosity most probably is stimulated differently. We here must capture how the curiosity that drives one to communicate (by sharing content) is impacted by the other members participating in the group. To that end, we propose several new metrics that help uncovering important components of (social) curiosity stimulation. As such, our work is completely orthogonal to [178].

5.3 WhatsApp

WhatsApp has become one of the main communication platforms in many countries [122]. It allows for one-to-one and group conversations, both encrypted. Groups are, by default, private spaces limited to 256 simultaneous users, although a user may join and leave a group at any time. Yet, as shown in [122], group administrators can make a group publicly accessible by sharing an invitation link in public websites, since anyone with the link can join the group. By gathering a large number of such publicly available invitation links, researchers were able to join the groups automatically and, once a member, gather data for posterior analysis.

For example, some researchers developed automatic tools to expose, in an anonymized fashion, the content being shared in publicly accessible groups [98, 40], whereas others analyzed properties of such content [106] with notable focus on pieces of information that had been previously checked as fake by fact checking agencies [122, 121, 94, 23]. Orthogonally, Melo *et al.* analyzed whether limiting message forwarding could mitigate misinformation

spread on WhatsApp [98]. They found that, though effective in slowing down the process, such approach would not stop misinformation from being widely distributed.

In a complementary directions, a few prior studies focused on exposing the importance of the underlying networks that connect users across different WhatsApp groups to information spread [122, 112]. In [122], the authors analyzed the structural properties of the network built from connecting users belonging to the same group, finding that this network has several properties, often observed in other online social networks, that had been associated with content virality. More recently, Nobre *et al.* focused on the media co-sharing network, built by connecting users who shared the same content, revealing the presence of strongly connected user communities that consistently help speeding up information spread [112].

Despite the recent surge of interest in analyzing content properties and user sharing patterns in WhatsApp, we are not aware of any prior attempt to investigate how to model social aspects of curiosity in this platform, as we do here. We believe that understanding the drivers behind users' sharing actions is key to proper modeling and understanding information dissemination on the platform. We take a step in that direction by proposing metrics to quantify one such driver – social curiosity.

5.4 Novel Metrics of User Curiosity

In this section, we present our novel metrics of curiosity stimulation driving user participation in group communication. Although we use WhatsApp as case study, the metrics are derived to be applicable to group communication in general. We consider that group membership is naturally dynamic, as members join or leave the group at their will. However, it is expected that (a subset of) members remain interacting with each other, exchanging opinions and content in general, for some time. We here are interested in quantifying the extent to which the bond created by such interactions can stimulate a member's curiosity towards carrying on the conversation. Specifically, we focus on quantifying stimuli to participating in such group conversations by *sharing content*. In other words, we propose metrics that capture different aspects of the stimuli one is driven by when choosing to share a piece of content with the group. As mentioned before, our derivation of the proposed metrics follows Berlyne's arguments and his proposed methodology [16]. Specifically we make use of measures from information theory to derive metrics that aim to capture different collative variables associated with curiosity stimulation.

In the following, we first present the notation used to define our metrics (Section 5.4.1). We then introduce our novel metrics of social curiosity, which aim to capture social

influence as a stimulus to curiosity and constitute a key contribution of this work (Section 5.4.2). Next, we present metrics that instantiate other traditionally studied collative variables to the particular context of WhatsApp groups (Section 5.4.3). The latter are inspired by the metrics introduced in [147], originally developed for the context of online music consumption, and here adapted to a very different domain, notably content sharing in a group communication platform. They are used in this work for comparison purposes, so as to emphasize the complementary role of social influence as a component of curiosity stimulation. We defer to Section 5.8 a discussion on limitations introduced by some of our assumptions and design decisions and their implications to the study.

5.4.1 Notation

We consider a universe of analysis consisting of sequences of messages shared by a number of users in various (independent) groups. A message is composed of pieces of content in any of four different media types, namely text, image, audio or video. Figure 3.3 of Chapter 3 illustrates the case of a sequence of 12 messages shared by 4 user members of a given group in the interval between 1:00pm and 3:00pm. The left part of the figure shows the sequence as seen by user 3, with the messages shared by her highlighted in blue. Note that the figure shows, for each message, the user who shared it, the message identifier, the content’s media type and the time of sharing. In the following, as in the next sections, we will use this example to illustrate the main concepts behind our metrics.

We now introduce the notation used to derive them in the following sections. We start by defining the sets of users, groups, and types of media used to compose the sequence of messages under analysis as $\mathcal{U}$, $\mathcal{G}$, $\mathcal{M}$, respectively. We do not assume access to message content, as it may not be available. Instead, we use the different types of media used to compose the message content as a representation of it, that is, we define $\mathcal{M} = \{\text{“text”}, \text{“image”}, \text{“audio”}, \text{“video”}, \text{“URL”}\}$ ². Each message under analysis is then a tuple $(u, \mathcal{C}_m, g, t)$ indicating that it was shared by user $u \in \mathcal{U}$ in group $g \in \mathcal{G}$ at time t . Set $\mathcal{C}_m \subseteq \mathcal{M}$ includes all media types used to compose the message’s content. We refer to such media types as the *categories* of the message. Note that the same message may contain content in different medias (e.g., a text and an image), thus a message may have multiple categories captured in set $\mathcal{C}_m$ ³.

We also note that, even though URLs are presented in textual format, we do con-

²Other possible categories include emojis, stickers and gifs. However, these are not available in our dataset, and thus are not explored in this work.

³The same media type used multiple times to build a message’s content (e.g., multiple images associated with it) is counted only once as part of $\mathcal{C}_m$.

sider them a separate category because URLs refer to external content (e.g., a post in another platform), which is not immediately visible to the user. Thus, from the perspective of curiosity stimulation, we speculate they may have a different impact compared to the rest of the textual content (which is immediately visible). More broadly, the use of media types as message categories is based on the assumption that different media types require different amounts of effort from the user to see and process the content, which should impact how the user curiosity is stimulated by it. We further elaborate on this point and discuss the implications of our choice of media types as categories in Section 5.8.

We use notation $t_{i|u,g}$ to refer to the timestamp of the i^{th} message shared by user u in group g . We assume time is continuous starting at 0. Thus, no two messages can be shared at exactly the same time in a group. Each group is composed by a set of users $\mathcal{U}_g \subseteq \mathcal{U}$, with size $|\mathcal{U}_g|$. Moreover, let the groups be associated with the random variable G .

These definitions illustrate the basis of our notation. Variables are represented as small letters (e.g., user u and group g), random variables as capital letters (e.g., G), and sets are represented as calligraphic letters (e.g., $\mathcal{G}$ and $\mathcal{U}$). We employ subscripts to determine subsets (e.g., users in a group $\mathcal{U}_g$), and the Bayesian notation to define constraints/dependencies (e.g., $t_{i|u,g}$ is the time of the i^{th} message shared by a given user u in a given group g). The complete set of notation, including those used above as others introduced in the following sections, is listed in Table 5.1.

5.4.2 Social Curiosity

As with the first case study of LastFM in Chapter 4, in this chapter we used the guidelines of [16] to create metrics that capture social influence as a stimulus for user curiosity. This means that we need to specify the *information spaces* in the context of WhatsApp and use the metrics of *information theory* to define our novel metrics. To this end, we recall that the collative variables capturing the stimulus are influenced by the type of stimulus that occurred and what *stimuli/responses* could have occurred and how likely they were [16]. Therefore, our metrics for capturing social influence also depend on the degree of uncertainty, i.e., Shannon entropy, which indicates the average amount of information in this *information spaces*.

In this case study, our goal is to analyse and model user curiosity. To this end, we consider social influence as a stimulus for curiosity in the context of a group communication platform such as WhatsApp. Therefore, the *information space* is defined as

Table 5.1: Main notation used to derive the curiosity stimulation metrics.

Notation	Description
$\mathcal{U}$	set of all users (individual user indicated by u)
$\mathcal{G}$	set of all users (one group indicated by g)
$\mathcal{M}$	set of all media types (i.e., categories) used as content representation ($\mathcal{M} = \{\text{"text"}, \text{"image"}, \text{"audio"}, \text{"video"}, \text{"URL"}\}$)
$\mathcal{C}_m$	set of categories associated with a given message (one category indicated by c)
$\mathcal{U}_g$	set of user members of group g
$t_{i u,g}$	timestamp of the i^{th} message shared by user u in group g
δ_T	duration of window of interaction
u	curious user whose curiosity we are analyzing
s	stimulating user who may stimulate the curiosity of a curious user at a given time
$\rightarrow$	indication of the quantity calculated for the current window of interaction
$\leftarrow$	indication of the quantity calculated for historical windows
$\mathcal{U}_{t,g}^{\rightarrow}$	set of (distinct) users sharing content in group g during window of interaction $[t - \delta_T, t]$
$\mathcal{S}_{t_{i u,g}}^{\rightarrow}$	set of (distinct) stimulating users for curious user u in group g at time $t_{i u,g}$ (all users, except for u , who shared content during the window $[t_{i u,g} - \delta_T, t_{i u,g}]$)
$\mathcal{C}_{t,g}^{\rightarrow}$	set of (distinct) message categories shared during window $[t - \delta_T, t]$
$\mathbf{S}_{t,g}^{\leftarrow}$	contingency table with historical patterns computed for group g at time t (see Figure 5.1)
$n_{u t,g}^{\rightarrow}$	number of times user u shared content in group g during window $[t - \delta_T, t]$
$n_{c t,g}^{\rightarrow}$	number of messages of category c shared in group g during window $[t - \delta_T, t]$
$n_{s,u t,g}^{\leftarrow}$	number of times s and u shared content in group g (in that order) within a δ_T time interval <i>before</i> the window $[t - \delta_T, t]$
G	random variable associated with groups
S	random variable associated with stimulating users
U	random variable associated with curious users
C	random variable associated with the categories of messages
$P_{t,g}^{\leftarrow}(X)$	Probability of random variable X computed based on historical patterns within group g before window of interaction $[t - \delta_T, t]$
$PMI_{t,g}^{\leftarrow}(U=u, S=s)$	Pointwise mutual information between the curious user u and stimulating user s in group g , measured at time t (Equation 5.1)
$socInf_{t,g}^{\leftarrow}(U=u, S=s)$	Social influence from stimulating user s on curious user u in group g , measured at time t (Equation 5.2)
$MI_{t,g}^{\leftarrow}(U, S=s)$	mutual information of all curious users conditioned on an particular stimulating user s in group g at time t (Equation 5.5)
$indSocInf_{t,g}^{\leftarrow}(U, S=s)$	indirect social influence from stimulating user s on all curious user (users in U) in group g , measured at time t (Equation 5.6)
$H_{t,g}^{\leftarrow}(U S=s)$	entropy of curious users conditioned on a particular stimulating users s in group g , measured at time t (Equation 5.9)
$groupEntropy(t=t_{i u,g}, g)$	group-level entropy of curious users conditioned on all stimulating users in group g , measured at time t (Equation 5.10)
$H_{t,g}^{\leftarrow}(U)$	entropy of curious users in group g , regardless of any social influence, measured at time t (Equation 5.11)

follows:

- (a) *Stimulus space*: is the set of shared messages by the *stimulating* users which is considered as *input signal*;
- (b) *Stimulus response*: it the set of shared messages by the *curious* users which is considered as *output signal*.

In these *information spaces*, the competing strengths of the possible *stimuli/responses* are the probabilities of the classes (e.g., *stimulating* users and *curious* users), as they are also given by [16]. However, the frequency of the individual possible alternatives of classes is indicated by the access history. This means that the historical interactions between users in WhatsApp is used to take into account the frequency of users' sharing behavior and thus calculate the probabilities of competing strengths (i.e., classes).

Having introduced the general concepts, notation and assumptions, we continue by taking a closer look at the novel metrics that capture the impact of social influence as the main driver of user curiosity (or simply social curiosity [63]). Recall that these metrics are meant to capture the influence that a set of users have on a particular (distinct) target user, as stimuli to this user's curiosity towards sharing content. Thus, in order to distinguish between the user whose curiosity we are analyzing, i.e., user u who shared a message in group g at time $t_{i|u,g}$, and the other users who may stimulate u 's curiosity through social influence at time $t_{i|u,g}$, we call the former a *stimulating* user and the latter a *curious* user (via social influence).

In the following, we use notations s to refer to a particular stimulating user and u to refer to one of his curious user. We also define $\mathcal{S}_{t_{i|u,g}}^{\rightarrow}$ as the *set* of all stimulating users for the curious user u sharing a message at time $t_{i|u,g}$, i.e., the set of all other users who shared content in the window of interactions $[t_{i|u,g} - \delta_T, t_{i|u,g}]$ and thus may stimulate u 's curiosity at time $t_{i|u,g}$. Note that, by definition, the curious user u herself does not belong to $\mathcal{S}_{t_{i|u,g}}^{\rightarrow}$, regardless if she also shared content during that period. In other words, we are interested in looking into the *other* users stimulating u 's curiosity. Also, as a set, multiple occurrences of the same stimulating user in the window of interaction count as one.

Given our assumption that influence is estimated based on historical patterns (assumption 5), we must distinguish between the *current* window of interaction and previous/historical windows (always with duration δ_T) which occurred before the current window initiated. Historical windows are used to estimate the influence of each stimulating user s in u 's content sharing at time $t_{i|u,g}$. We make such distinction by adopting notation $\rightarrow$ (as in $\mathcal{S}_{t_{i|u,g}}^{\rightarrow}$) to refer to the current window of interaction and $\leftarrow$ to refer to historical windows. Moreover, since we assume social influence may change over time, it must be recomputed using all available historical interactions up to (the beginning of) the current window of interaction.

		Curious users (target)				Total of rows:
		u_1	u_2	u_3	u_4	
users that stimulating the curiosity	u_1	0	1	2	1	4
	u_2	0	0	2	0	2
	u_3	1	0	0	1	2
	u_4	1	0	2	0	3
Total of columns:		2	1	6	2	11
						Total:

Figure 5.1: Contingency table $\mathbf{S}_{t,g}^{\leftarrow}$ computed when user 3 shares message 12 in Figure 3.3 of Chapter 3.

For each window of interaction defined by a message sharing at time $t = t_{i|u,g}$, historical interactions are captured by a contingency table $\mathbf{S}_{t,g}^{\leftarrow}$. Each cell in the table determines the number of times that stimulating user s and curious user u shared content in group g (in that order) within a maximum time interval of δ_T before the current window of interaction started (i.e., before time $t - \delta_T$), with $s, u \in \mathcal{U}_g$ and $s \neq u$. We use notation $n_{s,u|t,g}^{\leftarrow}$ to refer to this number. Note that, by definition, $n_{u,u|t,g}^{\leftarrow} = 0$, i.e., self-influence (diagonal) is not considered. For the general case of a group with $|\mathcal{U}_g|$ users, table $\mathbf{S}_{t,g}^{\leftarrow}$ has size $|\mathcal{U}_g| \times |\mathcal{U}_g|$. Now, it is possible to calculate the *competing strengths*, i.e., the probabilities, from the contingency table $\mathbf{S}_{t,g}^{\leftarrow}$.

Figure 5.1 shows the contingency table computed for the example in Figure 3.3 at 3:00pm, when user 3 shared *msg 12*. All messages shared up to the beginning of the window of interaction defined for *msg 12* should be considered. Thus, *msgs 1–8*, shared before 2:30, are taken into account to compute the table. The column representing user 3 as curious user is shown in red. For example, we can see that user 3 followed a message by user 1 within $\delta_T = 30$ minutes twice in the past (specifically when user 3 shared *msgs 3* and *8*). Similarly, user 3 followed messages by user 4 within that interval also twice in the past: *msg 8* by user 3 follows *msgs 5* and *7* by user 4. The figure also shows the joint count of the number of messages for each stimulating user (row), each curious user (column) and the overall total.

Now, let S and U be random variables associated with stimulating users (from *stimulus space*) and curious users (from *response space*), respectively. Given the contingency table computed for a message shared at time t by curious user u , we define the following probabilities (i.e., *competing strengths*) for a group g , computed based on historical patterns (note the use of $^{\leftarrow}$):

- (a) Probability of an stimulating user s on group g :

$$P_{t,g}^{\leftarrow}(S=s) = P^{\leftarrow}(S=s | t_{i|s,g}=t, G=g) = \frac{\sum_u n_{s,u|t,g}^{\leftarrow}}{\sum_{s,u} n_{s,u|t,g}^{\leftarrow}},$$

intuitively, this represents the likelihood that a stimulating user has shared a message and is considered as a social *stimulus* for the curious user;

- (b) Probability of a curious user u on group g :

$$P_{t,g}^{\leftarrow}(U=u) = P^{\leftarrow}(U=u | t_{i|u,g} = t, G = g) = \frac{\sum_s n_{s,u|t,g}^{\leftarrow}}{\sum_{s,u} n_{s,u|t,g}^{\leftarrow}},$$

in simple terms, this refers to the probability that a curious user has shared a message and is considered the curious user's *response* to the social stimuli received;

- (c) Probability of a curious user given an stimulating user:

$$P_{t,g}^{\leftarrow}(U=u | S=s) = P^{\leftarrow}(U=u | S=s, t_{i|u,g} = t, G=g) = \frac{n_{s,u|t,g}^{\leftarrow}}{\sum_u n_{s,u|t,g}^{\leftarrow}},$$

this represents the conditional probability of a curious user to share a message when the stimulating user has shared a message. Basically, it is considered as the possibility of a *response* from the curious user based on a previous *stimulus* from the stimulating user;

- (d) Joint probability of a curious user and an stimulating user:

$$P_{t,g}^{\leftarrow}(U=u, S=s) = P^{\leftarrow}(U=u, S=s, t_{i|u,g} = t, G=g) = \frac{n_{s,u|t,g}^{\leftarrow}}{\sum_{s,u} n_{s,u|t,g}^{\leftarrow}},$$

intuitively, the joint probability is the likelihood of a stimulating user sharing a message together with a curious user (i.e., *stimulus* and *response*) in the same window of interaction.

Note that all probabilities are conditioned on the group and timestamp of the current message. Thus, to simplify notation we drop these two conditions, using subscripts instead. That is, we define $P_{t,g}^{\leftarrow}(\cdot) = P^{\leftarrow}(\cdot | t_{i|u,g} = t, G = g)$ as exemplified above⁴. We make the same simplification in the notation used below to improve readability.

Taking the contingency table shown in Figure 5.1 as input, we can make, for example, the following computations at the time when user 3 shares *msg 12*:

(a) Probability of stimulating user 3: $P_{t,g}^{\leftarrow}(S=u_3) = \frac{2}{11};$

(b) Probability of curious user 3: $P_{t,g}^{\leftarrow}(U=u_3) = \frac{6}{11};$

- (c) Probability of curious user 3 given stimulating user 2:

$$P_{t,g}^{\leftarrow}(U=u_3 | S=u_2) = \frac{2}{2} = 1;$$

- (d) Probability of curious user 3 given stimulating user 4:

$$P_{t,g}^{\leftarrow}(U=u_3 | S=u_4) = \frac{2}{3};$$

⁴In other words, unless otherwise noted, whenever t is used we refer to time $t_{i|u,g}$ when the i^{th} message of the target curious user u was shared in group g .

(e) Joint probability of curious user 3 and stimulating user 4:

$$P_{t,g}^{\leftarrow}(U=u_3, S=u_4) = \frac{2}{11}.$$

With these probabilities in hand, we are ready to define our new metrics for social curiosity. In total, we propose four new metrics for the stimulation a user is exposed to by the social influence of others in the group. All metrics are based on the concept of *mutual information* [32]. In essence, the Mutual Information (MI) of two random variables is a measure of the mutual dependence between them. Specifically, it quantifies the reduction in the *uncertainty* (or Shannon entropy [92]) associated with one variable as we observe the other random variable.

Our first two metrics are based on the concept of PMI, which is calculated for a pair of outcomes x and y belonging to discrete random variables X and Y . In our case, we define the pointwise mutual information of a particular curious user u and a particular stimulating user s as:

$$PMI_{t,g}^{\leftarrow}(U=u, S=s) = \log_2 \left(\frac{P_{t,g}^{\leftarrow}(U=u | S=s)}{P_{t,g}^{\leftarrow}(U=u)} \right). \quad (5.1)$$

The PMI of u and s measures the *reduction* in the uncertainty of curious user u sharing a message due to the knowledge that s also shared a message (within the window of interaction).⁵ Intuitively, this metric captures how much knowledge about the social behavior of a particular stimulating user *reduces* the Shannon information content about the social behavior of the curious user. This *reduction* comes from the social influence that s has on u , which in turn stimulates u 's curiosity. It means that, based on historical patterns (social influence estimate), the behavior of (i.e., sharing of a message by) user u in the group is influenced by recent behavior of user s . Naturally, we set the PMI value to 0 if there is no history of messages shared by u , i.e., $P_{t,g}^{\leftarrow}(U=u) = 0$.

Note that, for the sake of estimating the social influence of s over u as a stimulus to u 's curiosity, we are interested in *positive values* of PMI. However, there are cases in which Equation 5.1 returns non-positive values. For example, $P_{t,g}^{\leftarrow}(U=u | S=s) = P_{t,g}^{\leftarrow}(U=u)$ suggests that the actions of u and s are independent, i.e., s has no influence of u 's behavior. In that case, the PMI value is equal to 0. Similarly, if $P_{t,g}^{\leftarrow}(U=u | S=s) = 0$, there is no historical evidence of messages shared by s and u within a maximum interval δ_T , and, as defined, there is no evidence of social influence from s on u . In addition, the PMI value will be negative whenever $P_{t,g}^{\leftarrow}(U=u | S=s) < P_{t,g}^{\leftarrow}(U=u)$. This case implies that the particular pair of curious user u and stimulating user s occurs less frequently than would be expected under the assumption of independent behavior. This means that although user s and user u have interacted with each other in the past, this interaction is not

⁵Alternative information theory-based metrics, such as transfer entropy [136, 164] could also be used. We here chose PMI, instead, as it refers to single events, i.e., a single message shared by a user. Transfer entropy, instead, being based on mutual information, is meant to capture an aggregate relationship between two processes – specifically, the amount of directed transfer of information between them.

sufficient to influence the curious user u .⁶ Regardless, in all such cases, there is clearly no evidence of social influence of s on u .

Taking into account the concepts on PMI values explained in the previous paragraph, we can summarise them as follows:

- (a) *positive PMI*: occurs when the frequency of interactions between s and u in the universe of messages shared by s is greater than the frequency of messages shared by u in the universe of messages shared within the group, i.e, $P_{t,g}^{\leftarrow}(U=u|S=s) > P_{t,g}^{\leftarrow}(U=u)$. Basically, this means that the *action events* (the sharing of messages) of the stimulating user s change the *action events* of the curious user u ;
- (b) *zero PMI*: takes place when the behavior of u is completely independent of s , although they have interacted with each other in the past according to historical patterns, i.e, $P_{t,g}^{\leftarrow}(U=u|S=s) = P_{t,g}^{\leftarrow}(U=u)$. Intuitively, the sharing actions of the stimulating user s do not change the sharing actions of the curious user u ;
- (c) *negative PMI*: occurs when the stimulating user s is not able to influence the curious user u because the frequency of interactions between s and u in the universe of messages shared by s is lower than the frequency of messages shared by u in the universe of messages shared within the group, i.e, $P_{t,g}^{\leftarrow}(U=u|S=s) < P_{t,g}^{\leftarrow}(U=u)$. In simple terms, the strength of social influence of the stimulating user s is lower than the participation strength of the curious user u in terms of sharing messages in the group.

In this context, as in prior studies [149], we choose to clip negative PMI values at 0 and leave the investigation of negative PMI values and their individual influences to future work. Thus, we define the following estimate of the social influence of the stimulating user s on the curious user u :

$$socInf_{t,g}^{\leftarrow}(U=u, S=s) = \max\left(PMI_{t,g}^{\leftarrow}(U=u, S=s), 0\right). \quad (5.2)$$

Given that the current window of interaction may have multiple stimulating users stimulating the curiosity of curious user u , we need to aggregate the mutual information computed for each stimulating user s that appears in set $\mathcal{S}_{t_i|u,g}^{\rightarrow}$. We do so by computing the average and maximum social influence for all stimulating users s on u . We thus build two metrics capturing the average and the maximum *direct influence* of the stimulating users on a curious user u , *based on historical patterns*:

$$avgDirInf(u, t=t_i|u,g, g) = \frac{\sum_{s \in \mathcal{S}_{t_i|u,g}^{\rightarrow}} socInf_{t,g}^{\leftarrow}(U=u, S=s)}{|\mathcal{S}_{t_i|u,g}^{\rightarrow}|}, \quad (5.3)$$

⁶In particular, a negative PMI helps to increase the amount of information needed to describe the event of curious user's social behavior given stimulating user's social behavior, i.e., $-\log_2 P(U=u|S=s)$. This means that the behavior of u cannot be explained by the social influence of s . Furthermore, in this dissertation we focus only on the positive PMI values, once they provide clear evidence of social influence, and leave the analysis of the negative PMI values to a future work.

$$\maxDirInf(u, t=t_{i|u,g}, g) = \max_{s \in \vec{\mathcal{S}}_{t_{i|u,g}}} \left(\text{socInf}_{t,g}^{\leftarrow}(U=u, S=s) \right). \quad (5.4)$$

In addition to direct influence, we argue that some users may have a natural ability to influence others. This is not a novel concept, neither in general nor in the particular case of WhatsApp, although we are the first to propose a metric to quantify it as a component of curiosity stimulation. Indeed, Guo *et al.* [51] made a distinction between direct and indirect influence, emphasizing that the latter may not have an obvious manifestation and may have gradual consequences on user behavior. Caetano *et al.* [22], in turn, identified the presence of (what they called) activists in public WhatsApp groups, i.e., users who often share content and, in doing so, keep the discussion going by driving others to also share content. This latter observation, in particular, motivated us to include metrics to capture such *indirect* influence: in a constrained (i.e., limited to 256 simultaneous users) and often focused space of communication, such as a WhatsApp group, it is reasonable to expect that these very active users (and activists) may indeed influence others much more often than it would occur in other unconstrained environments. That is, we consider that, even in the absence of prior experiences, i.e., $P_{t,g}^{\leftarrow}(U=u | S=s) = 0$, such influencers may still stimulate the curiosity of a user (e.g., a newcomer) in the group.

We identify these *indirect* influencers by searching for stimulating users s that tend to have *strong* influence towards some curious user u . Given this rationale, we propose two other metrics of social curiosity to capture the impact of such indirect influencers on curious user u at time $t_{i|u,g}$. These metrics are based on the *mutual information of the curious users conditioned on a particular stimulating user*, that is:

$$MI_{t,g}^{\leftarrow}(U, S=s) = \sum_{u' \in \mathcal{U}_g} P_{t,g}^{\leftarrow}(U=u', S=s) PMI_{t,g}^{\leftarrow}(U=u', S=s). \quad (5.5)$$

We note that, as done in Equation 5.2, in order to capture *indirect social influence*, we clip non-positive values of mutual information at 0 as a reflection of no social influence from stimulating user s on any curious user $u \in U$. We thus define:

$$\text{indSocInf}_{t,g}^{\leftarrow}(U, S=s) = \max\left(MI_{t,g}^{\leftarrow}(U, S=s), 0\right). \quad (5.6)$$

The random variable of interest in Equations 5.5 and 5.6 is U , conditioned on a particular stimulating user s . Smaller values of $\text{indSocInf}_{t,g}^{\leftarrow}(U, S=s)$ suggest weaker and no clear influence of s over any user $u \in U$. Take, for example, the case when $P_{t,g}^{\leftarrow}(U=u' | S=s)$ is roughly uniform for all curious users u' . Based on the values of $P_{t,g}^{\leftarrow}(U=u' | S=s)$, stimulating user s does not strongly influence any particular curious user u . The mutual information in this case decreases. In contrast, larger values of $\text{indSocInf}_{t,g}^{\leftarrow}(U, S=s)$ imply that the influence of stimulating user s tends to be concentrated and stronger on fewer curious users. These are the cases we are searching for. That is, the higher the value of $\text{indSocInf}_{t,g}^{\leftarrow}(U, S=s)$, computed based on historical patterns, the higher the

indirect influence that s may have over the target curious user u under analysis, i.e., the user whose curiosity stimulation is being quantified⁷.

Again, we aggregate the above metric for all stimulating users in the window of interaction via the average and maximum functions, referring to them as average and maximum *indirect* influence:

$$avgIndInf(u, t=t_{i|u,g}, g) = \frac{\sum_{s \in \mathcal{S}_{t_{i|u,g}}^{\rightarrow} indSocInf_{t,g}^{\leftarrow}(U, S=s)} indSocInf_{t,g}^{\leftarrow}(U, S=s)}{|\mathcal{S}_{t_{i|u,g}}^{\rightarrow}|}, \quad (5.7)$$

$$maxIndInf(u, t=t_{i|u,g}, g) = \max_{s \in \mathcal{S}_{t_{i|u,g}}^{\rightarrow}} \left(indSocInf_{t,g}^{\leftarrow}(U, S=s) \right). \quad (5.8)$$

As mentioned above, the indirect social influence metrics are particularly useful for capturing the influence that a stimulating user can have on a curious user with whom they have not yet interacted. If this is a recurring pattern, the direct influence metrics, calculated based on $P_{t,g}^{\leftarrow}(U=u | S=s)$ may increase over time.

Going back to the example in Figure 3.3, let's estimate the metrics of social curiosity applied to user 3 at the time $t=3:00\text{pm}$, when *msg 12* was shared. To that end, we will use the contingency table shown in Figure 5.1. The computation of all metrics are shown in Table 5.2. Note that users 1, 2 and 4 are stimulating users of influence for user 3 at this time (i.e., they shared content during the window of interaction). Using Equations 5.1–5.4, we compute the average and maximum direct social influence of these stimulating users on user 3 as $avgDirInf(u_3, t=3:00, g) = 0.39$ and $maxDirInf(u_3, t=3:00, g) = 0.87$, respectively. Also, using Equations 5.5–5.8, we compute average and maximum indirect social influence of the stimulating users on user 3 as $avgIndInf(u, t=3:00, g) = 0.59$ and $maxIndInf(u, t=3:00, g) = 0.87$, respectively.

As with the traditional collative variables for the first case study of LastFM in Section 4.2.2 of Chapter 4, we emphasize that there are some properties for two metrics of direct social influence and two metrics of indirect social influence based on PMI and partial MI, as follows:

Properties of social influence metrics: (1) *this measure is zero if the two specific events are independent, (2) it can be negative, and (3) it is directly related to the Shannon information content.*

The first two properties are directly related to the values of PMI (i.e., *positive, zero and negative value*), which are briefly explained after definition of Equation 5.1. They are explained in more detail in Appendix B and illustrated in Section B.1 of this appendix.

⁷We note that it may seem at first that Equations 5.5 and 5.6 lack this target curious user u . Yet, recall that, as the metrics in Equations 5.1 – 5.4, it is conditioned on the time $t = t_{i|u,g}$ that the target curious user u shared her i^{th} message in group g . This condition was omitted from the equations simply to improve readability.

Table 5.2: Computing the metrics of social curiosity for curious user 3 at time $t=3\text{pm}$ (reference: Figure 5.1).

Metric	Value
$PMI_{3:00,g}^{\leftarrow}(U=u_3, S=u_1)$	-0.13
$PMI_{3:00,g}^{\leftarrow}(U=u_3, S=u_2)$	0.87
$PMI_{3:00,g}^{\leftarrow}(U=u_3, S=u_4)$	0.29
$socInf_{3:00,g}^{\leftarrow}(U=u_3, S=u_1)$	0.00
$socInf_{3:00,g}^{\leftarrow}(U=u_3, S=u_2)$	0.87
$socInf_{3:00,g}^{\leftarrow}(U=u_3, S=u_4)$	0.29
$avgDirInf(u, t = t_{i u,g}, g)$	0.39
$maxDirInf(u, t = t_{i u,g}, g)$	0.87
$MI_{t,g}^{\leftarrow}(U, S=u_1)$	0.42
$MI_{t,g}^{\leftarrow}(U, S=u_2)$	0.87
$MI_{t,g}^{\leftarrow}(U, S=u_4)$	0.48
$indSocInf_{t,g}^{\leftarrow}(U, S=u_1)$	0.42
$indSocInf_{t,g}^{\leftarrow}(U, S=u_2)$	0.87
$indSocInf_{t,g}^{\leftarrow}(U, S=u_4)$	0.48
$avgIndInf(u, t = t_{i u,g}, g)$	0.59
$maxIndInf(u, t = t_{i u,g}, g)$	0.87

The third property, in turn, refers to Shannon information content of a particular curious user (response space), because the metric based on PMI or partial MI is the surprisal reduction (e.g., the amount of information) due to the knowledge of two particular events (the sharing of content) from *stimulus space*, i.e., stimulating user, and *response space*, i.e., curious user. Moreover, PMI can be obtained by the difference of Shannon information content and the *conditional* Shannon information content [91] as follows:

$$PMI_{t,g}^{\leftarrow}(U=u, S=s) = -\log_2 P_{t,g}^{\leftarrow}(U=u) - \left(-\log_2 P_{t,g}^{\leftarrow}(U=u|S=s) \right).$$

Before moving forward, we note that mutual information can also be used to quantify how social influence drives curiosity of a group of users. As such, it may offer an aggregate view of (social) curiosity stimulation in the ecosystem of a particular group g up to time $t_{i|u,g}$ when a member u shares a message, i.e., it captures the *collective behavior* of social influence. Towards defining such a metric, we start by introducing the entropy of curious users conditioned on a particular stimulating user s [92], given by:

$$H_{t,g}^{\leftarrow}(U|S=s) = -\sum_{u' \in \mathcal{U}_g} P_{t,g}^{\leftarrow}(U=u'|S=s) \log_2 \left(P_{t,g}^{\leftarrow}(U=u'|S=s) \right). \quad (5.9)$$

We then aggregate the conditional probability in Equation 5.9 over all stimulating users in the group to build a group-level metric :

$$groupEntropy(t=t_{i|u,g}, g) = \sum_{s \in \mathcal{U}_g} P_{t,g}^{\leftarrow}(S=s) H_{t,g}^{\leftarrow}(U|S=s). \quad (5.10)$$

The *groupEntropy* metric should be analyzed in light of the entropy of curious users U regardless of social influence, defined as follows:

$$H_{t,g}^{\leftarrow}(U) = -\sum_{u \in \mathcal{U}_g} P_{t,g}^{\leftarrow}(U=u) \log_2 \left(P_{t,g}^{\leftarrow}(U=u) \right). \quad (5.11)$$

Based on Equations 5.10 and 5.11, we introduce the group-level mutual information between curious users and stimulating users to capture the reduction in uncertainty associated with *curious users* due to the knowledge of the social influence from other users (*stimulating users*) in the group. Specifically, we define:

$$\text{groupMutInf}(t=t_{i|u,g}, g) = H_{t,g}^-(U) - \text{groupEntropy}(t=t_{i|u,g}, g). \quad (5.12)$$

The higher the mutual information, the greater the influence of stimulating users over the curious users in group g . This may happen because either there are some highly influential stimulating users or group curiosity is often stimulated by many group members.

Mutual information at group-level, in turn, denotes the actual amount of information by which the uncertainty about the *input* is reduced when the *output* signal is received. In other words, the actual amount of information by which the uncertainty about the messages shared by *curious users* (*output signals* related to the curious users' behavior) is reduced when the *stimulating users* (*input signals* related to the stimulating users' behavior) share messages.

Going back once again to our example, we also estimate the group-level social curiosity at the time user 3 shared *msg 12* as follows. We first compute the entropy of curious users $H_{3:00,g}^-(U) = 1.91$ and the group-level entropy of curious users conditioned on stimulating users $\text{groupEntropy}(t=t_{i|u,g}, g) = 1.39$. Based on these two values, we compute $\text{groupMutInf}(t=3:00, g) = 1.91 - 1.39 = 0.52$ which corresponds to only 27% of the original entropy of curious users in the group. Thus the knowledge of the social influence from the stimulating users causes a reduction of only 27% ($\frac{0.52}{1.91}$) on the entropy of curious users. We could say that, up to this time, the group is not very strongly stimulated by social influence.

Finally, we can describe the following properties for the metric of *collective social influence* at the group level based on mutual information:

Properties of group level metric based on mutual information: (1) *this measure is zero if the two random variables are independent*, (2) *it is non-negative* and (3) *it is related to Shannon entropy and conditional Shannon entropy*.

These properties are inseparable from the mutual information [91] and they are explored and exemplified in the Appendix B for a better understanding. And the third property is demonstrated by the Equation 5.12, which describes the definition of this metric.

A summary of all user/individual and group level metrics of social curiosity is presented in the top part of Table 5.3. Unless otherwise noted, all metrics should be interpreted as: *the larger the value, the greater the curiosity stimulation*.

Table 5.3: Metrics of curiosity stimulation for content sharing in WhatsApp groups.

Collative Variable	Metric	Definition
Social curiosity (individual)	$avgDirInf(u, t, g)$	Average direct influence of all stimulating users on curious user u at time t in group g (Equation 5.3)
Social curiosity (individual)	$maxDirInf(u, t, g)$	Maximum direct influence of any stimulating user on curious user u at time t in group g (Equation 5.4)
Social curiosity (individual)	$avgIndInf(u, t, g)$	Average indirect influence of all stimulating users on curious user u at time t in group g (Equation 5.7)
Social curiosity (individual)	$maxIndInf(u, t, g)$	Maximum indirect influence of any stimulating user on curious user u at time t in group g (Equation 5.8)
Social curiosity (group)	$groupMutInf(t, g)$	Mutual information between curious users and stimulating users in group g at time t (Equation 5.12: large values imply many strong curiosity stimulators in group)
Novelty	$userNovelty(u, g, t)$	Novelty associated with curious user u sharing message in group g at time t (Equation 5.13)
Novelty	$catNovelty(C_m, g, t)$	Novelty associated with set C_m of categories of the message shared by curious user u in group g at time t (Equation 5.14)
Uncertainty	$userUncertainty(\mathcal{U}_g, g, t)$	Uncertainty associated with set $\mathcal{U}_g$ of users in group g , measured at time t (Equation 5.15)
Uncertainty	$catUncertainty(C_{t,g}^{\rightarrow}, g, t)$	Uncertainty associated with set $C_{t,g}^{\rightarrow}$ of categories of messages in group g , measured at time t (Eq. 5.16)
Conflict	$userConflict(\mathcal{U}_{t,g}^{\rightarrow}, g, t)$	Conflict associated with set $\mathcal{U}_{t,g}^{\rightarrow}$ of users sharing content in group g , measured at time t (Equation 5.17)
Conflict	$catConflict(C_{t,g}^{\rightarrow}, g, t)$	Conflict associated with set of $C_{t,g}^{\rightarrow}$ categories of messages shared in group g , measured at time t (Eq. 5.18)
Complexity	$catComplex(C_{t,g}^{\rightarrow}, \mathcal{M}, g, t)$	Complexity associated with set $C_{t,g}^{\rightarrow}$ of message categories and set $\mathcal{M}$ of messages in group g , measured at time t (Equation 5.20)

5.4.3 Other Collative Variables

Having introduced our metric of social curiosity, let us recall that the *information space* consists of the *stimulus space*, which represents the set of messages shared by the stimulating users (*input signal*), and the *response space*, which in turn represents the set of messages shared by the curious users (*output signal*). However, the main difference lies in the way the competing strengths (i.e. the probabilities of the classes) are calculated for the traditional collative variables. Then, the frequency of each possible alternative of the classes for the stimulus of the collative variable is given by the *current window of interaction* instead of the historical patterns, as we will see as follows.

In this context, now we present 7 metrics that instantiate other collative variables,

notably novelty, uncertainty, conflict and complexity, priorly associated with curiosity stimulation. We build on metrics originally proposed to the context of online music consumption [147], adapting them here to the domain of content sharing in group communication. Based on assumption 3, we assume once again that user curiosity is triggered based only on what happens (*who* shares *what*) during the window of interaction⁸ (thus the use of notation $\rightarrow$ throughout this section). Therefore, the metrics are computed based solely on properties of the *user* who is sharing the content (i.e., the curious user) and of the content being shared, as captured by the media types (i.e., categories) associated with it. A summary of these metrics is presented in the bottom part of Table 5.3.

We here propose two metrics that instantiate the *novelty* component of the stimulus imposed on curious user u at time $t_{i|u,g}$. The first one – *userNovelty* – captures the novelty related to the user sharing the content, i.e., user u . The idea is that the experience of sharing content is less novel to users who share more often, which ultimately affects the curiosity driving the action. This claim is inspired by arguments by Loewenstein [83], who states that novelty is determinant of exploratory behavior in the individual’s curiosity, and by Kashdan [63, 61], who argues that social curiosity may also be stimulated by novelty. Thus, we propose user novelty as a means to capture possible effects that a novel experience may have as a factor driving a user to participate in a group by sharing content. To our knowledge, we are the first to propose a novelty metric related to users, but we believe that in the specific domain of investigation – group communication – this variable may be a relevant component of curiosity stimulation. Indeed, as we will show in Section 5.5, the proposed metric is indeed complementary to the others in many cases, which is an indication that, to some degree, the metric does capture new information.

The *userNovelty* metric is defined as the *surprisal* associated with curious user u :

$$userNovelty(u, g, t = t_{i|u,g}) = \begin{cases} -\log_2(P_{t,g}^{\rightarrow}(U=u)), & \text{if } P_{t,g}^{\rightarrow}(U=u) > 0. \\ -\log_2(1/|\mathcal{U}_{t,g}^{\rightarrow}|), & \text{otherwise.} \end{cases} \quad (5.13)$$

where $P_{t,g}^{\rightarrow}(U=u)$ is the probability of u sharing content in group g during the current window of interaction, being thus defined as $n_{u|t,g}^{\rightarrow} / \sum_{u \in \mathcal{U}_g} n_{u|t,g}^{\rightarrow}$. Note that $P_{t,g}^{\rightarrow}(U=u) = 0$ corresponds to the largest surprisal associated with curious user u . In such case, we set the surprisal to its largest value possible, which corresponds to a uniform distribution of curious users, i.e., $-\log_2(1/|\mathcal{U}_{t,g}^{\rightarrow}|)$, where $\mathcal{U}_{t,g}^{\rightarrow}$ is the *set* of distinct users sharing content in group g during the current window of interaction $[t - \delta_T, t]$.

Recall that we assume users may share messages with content in the following categories, defined based on media type: $\mathcal{M} = \{\text{“text”}, \text{“image”}, \text{“audio”}, \text{“video”}, \text{“URL”}\}$. Given a particular message shared at time t , let C be the random variable associated with

⁸This was also a guideline for deriving the social curiosity metrics. However, for those metrics in particular, we used historical patterns to estimate the social influence of curious users and stimulating users identified in the window of interaction.

category, c be a category associated with the given message (i.e., $c \in \mathcal{C}_m$) and $\mathcal{C}_{t,g}^\rightarrow$ be the set of media types associated with all messages shared within the window of interaction $[t - \delta_T, t]$. Similarly to what was done for users, we define a novelty metric related to the message categories, based on the surprisal associated with that variable:

$$catNovelty(\mathcal{C}_m, g, t = t_{i|d,g}) = \begin{cases} -\log_2(\bar{P}_{t,g}^\rightarrow(\mathcal{C}_m)), & \text{if } \bar{P}_{t,g}^\rightarrow(\mathcal{C}_m) > 0. \\ -\log_2(1/|\mathcal{C}_{t,g}^\rightarrow|), & \text{otherwise.} \end{cases} \quad (5.14)$$

where $\bar{P}_{t,g}^\rightarrow(\mathcal{C}_m)$ is the average probability of categories associated with the message (categories in $\mathcal{C}_m$), defined as:

$$\bar{P}_{t,g}^\rightarrow(\mathcal{C}_m) = \frac{1}{|\mathcal{C}_m|} \sum_{c \in \mathcal{C}_m} P_{t,g}^\rightarrow(C=c).$$

The probability of a specific category is estimated as $P_{t,g}^\rightarrow(C=c) = n_{c|t,g}^\rightarrow / \sum_{k \in \mathcal{C}} n_{k|t,g}^\rightarrow$, where $n_{c|t,g}^\rightarrow$ is the number of messages with category c shared in group g in the window of interaction $[t - \delta_T, t]$ ⁹. This metric captures how surprising it is that the message shared at time $t = t_{i|d,g}$ has category c .

Taking the example of user 3 sharing *msg 12* at 3:00pm, illustrated in Figure 3.3, we find that a total of four messages were shared during the current window of interaction, defined as [2:30, 3:00]. The curious user, user 3, shared only once. Thus, the novelty associated with the curious user is computed as: $P_{3:00,g}^\rightarrow(U = u_3) = 1/4$, hence $userNovelty(u_3, t = 3:00, g) = 2$. Similarly, during this window, two images, one audio and one video were shared. In particular, *msg 12* contains only an image (i.e., $|\mathcal{C}_m| = 1$). Thus, the novelty associated with the category of this message can be computed as: $\bar{P}_{3:00,g}^\rightarrow(\mathcal{C}_m = \{\text{"image"}\}) = 2/4$, which results in $catNovelty(\{\text{"image"}\}, t=3:00, g) = 1$. Thus, the curiosity driving user 3 to share *msg 12* is more stimulated by the novelty the user herself experiences with this action than by the novelty of (the category of) the content shared.

Note that both novelty related metrics focus on a single element: the curious user or the (categories of the) message shared. The aggregation of either metric for all elements in the window of interaction, using the Shannon entropy, captures the *uncertainty*. Thus, we define metrics of uncertainty related to user and categories as:

$$userUncertainty(\mathcal{U}_g, g, t = t_{i|u,g}) = -\sum_{u \in \mathcal{U}_g} P_{t,g}^\rightarrow(U=u) \log_2(P_{t,g}^\rightarrow(U=u)), \quad (5.15)$$

$$catUncertainty(\mathcal{C}_{t,g}^\rightarrow, g, t = t_{i|u,g}) = -\sum_{c \in \mathcal{C}_{t,g}^\rightarrow} P_{t,g}^\rightarrow(C=c) \log_2(P_{t,g}^\rightarrow(C=c)). \quad (5.16)$$

⁹Note that since messages may have multiple categories, the same message may be counted multiple times in the numerator and in the denominator of $P_{t,g}^\rightarrow(C=c)$.

The idea behind these metrics is that the curiosity of the curious user may be more/less stimulated by the greater/less diversity in the users sharing messages (and in the categories of these messages) during the window of interaction.

For the sake of illustration, let's compute the stimulus related to user uncertainty experienced by user 3 when sharing *msg 12*. We first calculate the probability of each user sharing a message in the current window of interaction, which is equal to $1/4$ for all four users in the window. The user uncertainty is then computed as $userUncertainty(t = 3:00, g) = 2$. Turning to the stimulus related to the uncertainty associated with message categories, we compute probabilities $P_{3:00,g}^{\rightarrow}(C = \text{"video"}) = 1/4$, $P_{3:00,g}^{\rightarrow}(C = \text{"audio"}) = 1/4$ and $P_{3:00,g}^{\rightarrow}(C = \text{"image"}) = 2/4$, which leads to an uncertainty equal to $catUncertainty(t = 3:00, g) = 1.5$. Thus, at time 3:00 pm, the diversity (uncertainty) in users sharing content is greater than the diversity in content category. Thus, user 3's curiosity is more stimulated by the former.

As in the context of LastFM's online music consumption in Chapter 4, the properties of metrics for collative variables based on Shannon *entropy*, namely *uncertainty of users* of Equation 5.15 and *uncertainty of message categories* of Equation 5.16, are also applied in the WhatsApp context. So, we remember the following properties:

Properties of metric based on entropy: (1) *this measure reaches the maximum when all probabilities (i.e., competing strength) are equal, and (2) when they are equal, it increases with the number of alternative classes to which the stimulus could belong.*

The Section A.2 of Appendix A shows how these properties are explored in details by definition of axioms related to them. Furthermore, these properties are related to the probabilities of the classes (the distinct alternatives of stimulus) in the window of interaction. So, when the probabilities of competing strengths are equal and more classes exist, the greater the uncertainty of the users or message categories within the window of interaction.

According to Berlyne [16], *conflict* occurs when the same stimulus triggers multiple incompatible responses, being positively related to the strengths of the competing responses. To operationalize the concept of conflict, we follow the same approach in [171]. The basic idea is that the different elements (categories/users) that appear in the window of interaction represent the potentially incompatible responses stimulating the curiosity of the curious user, and the strength of each response is captured by the probability of occurrence of each element (category/user). We instantiate this variable by first computing the average probability of users (categories) over all users (message categories) in the window of interaction, and then taking the surprisal of the result. This procedure leads to two new metrics, one related to users and the other to categories:

$$userConflict(\mathcal{U}_{t,g}^{\rightarrow}, t = t_{i|u,g}) = -\log_2(\bar{P}_t^{\rightarrow}(\mathcal{U}_{t,g}^{\rightarrow})), \quad (5.17)$$

where $\bar{P}_t^{\rightarrow}(\mathcal{U}_{t,g}^{\rightarrow})$ is the average probability of users calculated over all users in the interaction window defined by

$$\bar{P}_t^{\rightarrow}(\mathcal{U}_{t,g}^{\rightarrow}) = \frac{1}{|\mathcal{U}_{t,g}^{\rightarrow}|} \sum_{u \in \mathcal{U}_{t,g}^{\rightarrow}} P_{t,g}^{\rightarrow}(U=u),$$

and

$$catConflict(\mathcal{C}_{t,g}^{\rightarrow}, g, t=t_{i|u,g}) = -\log_2 \left(\bar{P}_t^{\rightarrow}(\mathcal{C}_{t,g}^{\rightarrow}) \right). \quad (5.18)$$

where $\bar{P}_t^{\rightarrow}(\mathcal{C}_{t,g}^{\rightarrow})$ is the average probability of categories calculated over all message categories in the interaction window defined by

$$\bar{P}_t^{\rightarrow}(\mathcal{C}_{t,g}^{\rightarrow}) = \frac{1}{|\mathcal{C}_{t,g}^{\rightarrow}|} \sum_{c \in \mathcal{C}_{t,g}^{\rightarrow}} P_{t,g}^{\rightarrow}(C=c). \quad (5.19)$$

Again, we use the example in Figure 3.3 to exemplify how the metrics of conflict are computed for user 3 sharing *msg 12*. We first take the average probabilities across all users and categories which are 1/4 and 1/3, respectively. We then compute $userConflict(t = 3:00, g) = 2$ and $catConflict(t = 3:00, g) = 1.59$.

Finally, an alternative form of capturing the diversity of message categories in the current window of interaction is by exploiting the unique occurrences of the media types. This metric, also based on surprisal, captures the *complexity* associated with the categories of messages shared during the current window of interaction, which is also a collative variable associated with curiosity stimulation. It is defined as:

$$catComplex(\mathcal{C}_{t,g}^{\rightarrow}, \mathcal{M}, g, t = t_{i|u,g}) = -\log_2 \left(\frac{|\mathcal{C}_{t,g}^{\rightarrow}|}{|\mathcal{M}|} \right). \quad (5.20)$$

Note that, unlike uncertainty and conflict, which captures the diversity of message categories considering only those included in the window of interaction, complexity captures a somewhat different notion of diversity that takes into account all possible categories (included in set $\mathcal{M}$). Once again, we compute the complexity of message categories at time 3:00 when user 3 shared *msg 12* as follows. The current window of interaction includes 3 types of media (notably, “video”, “audio” and “image”) out of a total of 5 different types possible. Thus, the complexity associated with message categories is computed as $catComplex(t=3:00, g) = -\log_2 3/5 = 0.74$.

In addition to the properties of the metrics for collative variables based on Shannon *entropy*, the metrics for collative variables based on Shannon *information content* also have some properties, as in the first case study of Chapter 4. So we recall these properties, but now in the WhatsApp context as follows:

Properties of metrics based on surprisal: (1) the lower this probability (i.e., competing strength) is, the more information the stimulus contains, and (2) no stimulus within the interaction window can have a higher surprisal than a stimulus that is not within the window of interaction.

The first property applies to metrics of *novelty* (Equations 5.13 and 5.14), *conflict* (Equations 5.17 and 5.18) and *complexity* (Equation 5.20). The closer the probability of a class is to 1 (zero), the lower (higher) the value of these metrics and the less (more) information is available from the respective stimulus. The second property, on the other hand, refers to the novelty metrics such as the *novelty of user* (Equation 5.13) and the *novelty of message categories* (Equation 5.14). Thus, if the current message category (or curious user) is not within the interaction window, Shannon information content in this case is greater than for all message categories (or users) within the window of interaction. To better understanding these properties, we use some axioms defined in the Appendix A, more specifically, the Section A.2 tackle some examples to show how these properties are explored.

As a final remark, we note that the metrics defined in Equations 5.13–5.20 are always non-negative and should be interpreted as: the greater the value the higher the stimulus to the curiosity of the curious user. Having introduced the novel metrics of curiosity stimulation, we applied them to study curiosity stimulation behind content sharing in Whatsapp groups. We discuss the main results from our investigation next.

5.5 Curiosity Stimulation in WhatsApp Groups

In this section, we use the metrics introduced in the previous section to characterize user curiosity in WhatsApp groups. Our goal is to both show the applicability of the metrics and uncover behavioral traits in a widely used communication platform.

To that end, we rely on a dataset gathered by the WhatsApp Monitor¹⁰ [98], a system for collecting shared messages in publicly accessible WhatsApp groups. The dataset consists of a sequence of messages posted in a number of publicly accessible groups in Brazil over a period of great social and political turmoil in the country (April 1st 2018 to April 30th 2019) which includes the 2018 general elections in the country. The monitored groups are themed around political topics.

Each entry in the dataset consists of posting time, anonymized user identifier, group identifier and message categories (media types)¹¹. We note that even though WhatsApp currently allows users to reply to a particular message, thus creating an explicit link between messages, this feature was not available at the time our dataset was collected¹².

¹⁰WhatsApp Monitor is available online at <http://www.monitor-de-whatsapp.dcc.ufmg.br/>.

¹¹All user identifiers are indeed anonymized phone numbers as we are not able to identify multiple phone numbers belonging to the same person. Moreover message content, though collected by WhatsApp Monitor, cannot be made publicly available and, as such, was disregarded.

¹²Yet, our metrics might still be used (with some adjustments) when replies are available. In that

As such, any possible link must be inferred from the available data, notably from the time when the messages were shared, since message content is not accessible.

To avoid issues with data sparsity, for each group, we only consider users who shared at least 30 messages during the whole period of $\delta_T = 30$ minutes, and we only compute the metrics for sharing events with at least 10 messages in the window of interaction¹³. These lower bounds were imposed so as to be able to compute reliable values for the metrics, following the arguments in prior work [165] that state that samples with fewer than 10 elements often lead to estimations with low statistical power. We also ran some preliminary experiments with other bounds and based on observations from these experiments as well as our own prior experience with curiosity stimulation on LastFM [147], we found that fewer data points often lead to unreliable measures. After applying these filters, the resulting dataset, which is used in our study, is composed on 2,054,302 messages posted by 7,584 distinct users in 335 groups.

5.5.1 Relationships among Collative Variables

We start by quantifying the relationships among different collative variables, namely social influence, novelty, uncertainty, conflict and complexity, as captured by the proposed metrics. We aim at assessing the extent to which different metrics, notably the four novel metrics related to (individual-level) social influence (Equations 5.3, 5.4, 5.7 and 5.8) are able to capture aspects related to curiosity stimulation which are not covered by the other (more traditional) variables (Equations 5.13–5.20). To that end, we classify the metrics into *redundant* or *complementary* as to whether their effect on curiosity stimulation is mostly captured by the others (thus being redundant) or not.

Specifically, we first compute, for each user u in a group g , all metrics (11 in total) for each message shared by u in g . We then use these values to compute the Spearman correlation coefficient between each pair of metrics, for each (u, g) pair. Next, as in [147], we employ a heuristic to identify redundant metrics: two metrics are considered redundant for a given (u, g) pair if their correlation falls in the $(-1, -0.5)$ or $(+0.5, +1)$ ranges, which can be considered moderate-to-strong correlations. Otherwise, the correlation is deemed weak and the metrics are taken as *complementary*.

case, we argue that, even though the original message was the primary source of curiosity stimulation on the user sharing the reply, other messages shared recently could also compose the stimulus driving the user behavior, perhaps with a lower intensity. Thus, a weighting scheme could be employed to combine all signals, favoring the message being replied. Exploring such scheme is an interesting subject for future work.

¹³For the sake of research reproducibility, we make our anonymized data as well as our code publicly available at https://zenodo.org/record/5790153#.Yb00sb_MJE4.




Figure 5.2: Fraction of redundant (user, group) cases for each pair of metrics.

Figure 5.2 shows a matrix with the fractions of cases, i.e., (user, group) pairs, for which each pair of metrics was considered redundant. One key result is that the four social influence metrics, shown in the last four rows (and columns), are *complementary* when compared to the metrics related to the other collative variables for most (more than 77%) of the cases. Thus, these metrics indeed capture relevant aspects of curiosity stimulation which are not covered by traditionally studied metrics. We also observe that direct and indirect social influence do often capture distinct effects, as their related metrics cannot be considered redundant for a large fraction of the cases (67% for the metrics of maximum influence). Yet, given the same general effect (i.e., direct or indirect social influence), taking either the average or maximum across all stimulating users of influence is quite similar, as the corresponding metrics are redundant with respect to each other in most cases.

We can also note great redundancy between several traditional metrics, such as between uncertainty and novelty and between uncertainty, conflict and complexity, all related to message categories, as well as between uncertainty and conflict related to users. These redundancies are consistent with those observed in [147] for curiosity in online music consumption, and reflect the similar effects captured by those related metrics.

In the following, given our goal to study social influence as a component of curiosity stimulation, we focus our analyses on the two social influence metrics – maximum direct




Figure 5.3: Diversity of social curiosity stimuli across message sharings and users.

influence and maximum indirect influence – identified as complementary to each other in most cases.

5.5.2 Diversity and Dynamics of Social Curiosity

As a first analysis, we focus on how diverse user curiosity stimulation is in terms of the two complementary metrics of social curiosity. Recall that both metrics are computed for each *message sharing* event, performed by a user u in a group g at time $t_{i|u,g}$. Figure 5.3 shows results of the two metrics for individual messages (Figure 5.3a) as well as averages across all messages shared by the same user u (Figure 5.3b), both computed for specific groups (i.e., the same user in different groups appear separately). Each graph in the figure shows the number of elements (messages or users) with specific values of *maximum direct influence* (x -axis) and *maximum indirect influence* (y -axis). For both graphs, we zoomed in a specific region of the graph and inset it to the graph. We highlight that the regions zoomed in correspond to 97% and 92% of all messages and users, respectively. The graphs are meant to contrast the values of both social curiosity metrics for each message and user. As such, they also show the ranges of the values observed for both metrics considering individual messages as well as aggregated across all messages shared by the same user.

The graphs reveal great diversity in values of both metrics across messages and users. On one hand, there is a major concentration in smaller values for both metrics, highlighted in the insets. Smaller values suggest that, based on the historical patterns, the behavior of the curious user (in terms of content sharing actions) is only weakly influenced (if influenced at all) by how the others in the current window of interaction (stimulating users) behave. That is, the small amount of information (in terms of bits) about the interactions between these users suggest weaker social influence. In contrast, a fraction

of the messages as well as a fraction of the users exhibit very large values of both metrics (two to four times larger than the majority), suggesting that, in those cases, the behavior of the curious user is strongly influenced by past behaviors of the stimulating users. As a consequence, this produces more information (in terms of bits) about the interactions between those users, offering stronger evidence of social influence between them. Also, as the axes of both graphs show, direct influence is clearly spread over larger values, as a result of how the metric is computed (i.e., based on pointwise mutual information).

Moreover, by comparing the spreading of values (especially in the x -axis) on both graphs, we can infer that the diversity is great even if we look at the same user over time, i.e., over all the messages the user shared. In other words, the (social) curiosity of a user is stimulated quite differently for different users and even for the same user, considering different messages shared by her (potentially in different groups), which is consistent with some of our key assumptions presented in Section 3.2 of Chapter 3.

Given these observations, we delve deeper into the analysis of the diversity and dynamics of social curiosity by considering three levels of aggregation. First, we focus on individual messages shared by users in the groups (Section 5.5.2.1). Next, we consider all messages shared by the same user in a given group (Section 5.5.2.2). Finally, we look at social curiosity at the group level (Section 5.5.2.3). We discuss our results next.

5.5.2.1 Social Curiosity at the Message Level

Aiming at identifying common patterns of social stimulation driving the sharing of a message, we clustered the message sharing events based on the values of both metrics – $maxDirInf$ and $maxIndInf$. To that end, we employed the Mini-Batch K-means algorithm [138], a widely used parallel version of K-means adequate for large datasets, to cluster the values of curiosity stimulus (computed by the two metrics) associated with the messages shared by all users (over 2 million messages in total). We varied the number of clusters k from 2 to 18 and, for each value, we repeated each experiment 10 times and computed the Silhouette index [8] of the clustered results. This is a measure of how similar an element is to its own cluster (cohesion) compared to other clusters (separation). Larger values indicate better matching of elements into clusters. Using the results of Silhouette index, we set $k = 3$ as the best number of clusters, as it led to the highest Silhouette index among all values analyzed (equal to 0.55, which suggests that a reasonable clustering structure was obtained [64]).

Table 5.4 shows a description of the three identified clusters in terms of average values (and associated 95% confidence intervals) of both metrics as well as fraction of

Table 5.4: Message-level social curiosity stimulation profiles.

Cluster	% <i>msgs</i>	Max. Direct Influence (average $\pm 95\%$ C.I.)	Max. Indirect Influence (average $\pm 95\%$ C.I.)
Independent	72.6%	0.26 ± 0.000750	0.04 ± 0.000044
Indirect	14.4%	0.43 ± 0.002320	0.15 ± 0.000210
Dependent	13.0%	3.52 ± 0.006623	0.05 ± 0.000099

messages in each cluster. Note that the confidence intervals are quite tight, implying that the average values of both metrics are very representative of the elements in each cluster. By comparing these values for each metric¹⁴, we label the clusters as:

- (Socially) Independent: both metrics tend to assume the smallest values among the three clusters, implying less frequent co-occurrence of curious user and stimulating users, and suggesting that social influence plays a less important role in curiosity stimulation. This occurs in 73% of the messages;
- Indirect (Social Influence): compared to the other clusters, the maximum direct social influence tends to assume moderate values, whereas the maximum indirect influence assumes the largest values of the three clusters. Thus, in comparison to the other two clusters, indirect social influence seems to have a particularly important role on the curiosity stimulation in this case. This occurs in 14% of the messages;
- (Socially) Dependent: this cluster exhibits, by far, the largest values of direct social influence, suggesting more frequent co-occurrences of curious user and stimulating users, and thus much stronger social stimuli as captured by this metric. This happens in 13% of the messages.

5.5.2.2 Social Curiosity at the User Level

Having identified patterns of social stimulation associated with individual messages, we now use them to cluster users in terms of their social stimulation profiles. Given our assumption, supported by the results in Figure 5.3b, that social curiosity stimulation may vary across time for the same user, our goal is to build a user profile representation that reflects such dynamics using the patterns of social stimulation at specific messages as building blocks. We do so by first representing each user u by the sequence of social stimulation patterns (independent, indirect or dependent) associated with the sequence of messages shared by u . We then model such sequence using a User Behavior Model

¹⁴Note that we do not compare values of different metrics as they cover very different intervals. Rather, we analyze differences across the three clusters by looking at the values of each metric at a time.




Figure 5.4: Example UBMG for a user u with a sequence of message-level stimulation pattern equal to PPIDDDIPPIPIIIDDDIPP in a group g (P for independent, I for indirect and D for dependent).

Graph (UBMG) [99]. A UBMG is a graph where each node represents a pattern of social stimulation (state) and edges denote the transitions between patterns from a message to the next one, by the same user. Edge weights are the probabilities of such transitions. One UBMG is used to model a user in *a given group*. Thus, the same user in different groups are represented by multiple (possibly distinct) UBMGs, one for each (user, group) pair.

We illustrate the UBMG representation by considering a fictitious user u who shared 20 messages in a group g . Given the social stimulation pattern associated with each message shared by u in g – P for independent, I for indirect and D for dependent – we first build a representation of (u, g) as a temporally-ordered sequence of stimulation patterns. Let's say such representation is PPIDDDIPPIPIIIDDDIPP. That is, u 's curiosity stimulation in g is associated with the independent (P) pattern in the first two messages, then changes to the indirect (I) pattern in the third message, changing once again to the dependent (D) pattern, and so on. The corresponding UBMG is shown in Figure 5.4. Note that once in the *independent* state, user u has 50% of chance of remaining in this state in the next message and 50% of chance of changing to *indirect*. Once in the *indirect* state, u either goes back to the *independent* state, with 50% of chance, changes to *dependent*, with 33% of chance, or remains in the same state (17% of chance). Having built the UBMGs for all (user, group) pairs, we then clustered these UBMGs to uncover patterns of user curiosity stimulation. To that end, we employed once again the Mini-Batch K-Means algorithm, using the Silhouette score to select the best number k of clusters. Figure 5.5 shows the average Silhouette score (computed for 10 runs) along with 95% confidence intervals as a function of k . Note that the average Silhouette increases greatly with k until reaching a rough plateau for $k \geq 13$. Also, the wider confidence intervals for small values k indicate higher variability of the results. Yet, the intervals become quite tight for larger values of k .

Based solely on the Silhouette index results, one would choose a value of k around 13 (or just above this mark) as the number of clusters. However, such large number of clusters makes it difficult to identify distinguishing characteristics. Indeed, by manually




Figure 5.5: Silhouette score: determining the number of UBMG clusters.

analyzing the centroids of these clusters, we found great similarities among many of them. Thus, in the interest of obtaining a more interpretable set of UBMG profiles, along with a manual analysis of the centroids of all $k = 13$ clusters, we chose to set k equal to 5 clusters, as it delivers the best trade-off between clustering quality (i.e., Silhouette score) and interpretability of the properties characterizing each identified cluster¹⁵. On one side, the Silhouette score for $k = 5$ is 0.30, which, despite lower, is still indicative of an existing clustering structure [64]. On the other side, we observed that larger values of k led to mostly variants of the same clusters identified with $k = 5$.




Figure 5.6: Results from Principal Component Analysis applied to $k = 5$, 9 and 13 clusters.

We illustrate the latter point by applying Principal Component Analysis (PCA) to identify the three most important components for three different values of k , notably $k = 5$, $k = 9$ and $k = 13$. Results are shown in Figure 5.6. By comparing Figures 5.6a and 5.6b we can see that increasing the number of clusters from $k = 5$ to $k = 9$ simply led to splitting some of the clusters into multiple ones, but those are quite similar to each other. Note for example cluster A_1 , identified in black in Figure 5.6a. It is basically split into two closely related variants, clusters B_4 and B_6 shown in pink and purple, respectively, in Figure 5.6b. Similarly, by comparing results for $k = 9$ and $k = 13$, we note that cluster

¹⁵We found that restricting to 5 clusters facilitates interpretation of the results as the identified clusters are more clearly distinct and refer to different patterns.

Table 5.5: User-level social curiosity stimulation profiles: percentage of messages in each UBMG state.

UBMG state	User-level profiles				
	U_0	U_1	U_2	U_3	U_4
Independent	65.73%	87.72%	80.17%	70.25%	8.66%
Indirect	20.88%	6.23%	19.80%	0.20%	84.49%
Dependent	13.39%	6.06%	–	29.54%	6.85%

B_7 (in light pink in Figure 5.6b) is mostly split into two variants, clusters C_7 and C_9 (cyan and light pink, respectively) in Figure 5.6c.

Figure 5.7 shows the five user-level profiles, referred to as $U_0 \dots U_4$, identified by the centroids of the five UBMG clusters obtained. For each profile, the figure also shows the number of (user, group) pairs in the cluster, and the average number of messages per such pair. To help interpreting the UBMGs, Table 5.5 shows, for each profile, the average fraction of all messages sent by the user in the given group characterized by each message-level curiosity stimulation profile.

As shown in Figure 5.7a, profile U_0 , including almost 30% of all (user, group) pairs in our dataset, is characterized by a general trend of repeatedly staying in the same state of social stimulation (strong self-transitions). In other words, users in this profile tend to experience a rather stable curiosity stimulation process, being consistently driven by one of the three patterns over time, with occasional changes to a different pattern. In particular, as shown in Table 5.5, this user profile is dominated by the *independent* state, which characterizes over 65% of the messages, on average. In other words, users with this profile, who are the most active ones in terms of the average number of messages shared, tend to be mostly insensitive to social influence around two thirds of the time. Yet, social influence does take a more important role as a component of curiosity stimulation, either directly or indirectly, in roughly 35% of the messages shared by those users.

Users in profile U_1 , shown in Figure 5.7b, may also experience the three states, but changes towards the independent state are much more frequent, which, in turn, tends to dominate the overall profile. In other words, compared to U_0 , users in U_1 tend to have less diversity in terms of social curiosity stimulation patterns, suffering little impact from social influence most of the time (87% of the messages). They also tend to be much less active than users in U_0 . This cluster includes 11% of all (user, group) pairs.

Profile U_2 , shown in Figure 5.7c, is characterized by a mixture of independent and indirect states, though being also greatly dominated by the former (80% of the messages, as shown in Table 5.5). Unlike users in the other clusters, users in U_2 do not experience the dependent state. In other words, the curiosity stimulation of these users is mostly insensitive to *direct* social influence. Moreover, the strong self-transitions imply that users often remain repeatedly in the same state of curiosity stimulation. Around 31% of the (user, group) pairs in our dataset fall into this cluster.






(a) U_0 : 2,476 elements (383 msgs/user).(b) U_1 : 889 elements (90 msgs/user).(c) U_2 : 2,550 elements (156 msgs/user).(d) U_3 : 2,156 elements (283 msgs/user).(e) U_4 : 235 elements (67 msgs/user).

Figure 5.7: User-level social curiosity stimulation profiles: state transition diagrams (UB-MGs), each element is a (user, group) pair.

Profile U_3 , shown in Figure 5.7d, exhibits a distinguishing characteristic compared to the others. There are no transitions into the indirect state from the other states (only the self-transition). Users in this cluster either start in the indirect state and possibly move towards the other states or never experience it. Compared to users in U_0 , U_1 and U_2 , users in U_3 tend to fall in the dependent state more often (in roughly 30% of the messages), being thus more sensitive to direct social influence as a driver to curiosity stimulation. These users tend to be very active, falling behind only users in U_0 . This cluster includes almost one third of all (user, group) pairs (26%).

Finally profile U_4 , shown in Figure 5.7e, shows the distinguishing characteristic of much stronger transitions (including self-transition) to the indirect state, leading to profiles that are dominated by this state, which characterizes 85% of the messages, on average. The independent state, in turn, is much less prevalent compared to the other profiles, as can be noted by the much lower self-transition probability. Indeed, as shown in Table 5.5, the independent state happens in only 8% of the messages shared by users with this profile, on average. These are users whose curiosity stimulation tends to be often influenced by social factors but by *indirect* means, possibly because the lack of frequent activity prevents the formation of strong social ties that could influence user curiosity. Indeed users in U_4 , which account for only 3% of all (user, group) pairs in our dataset,




Figure 5.8: Time series of both metrics of social curiosity for different (user, group) pairs.

are the least active ones in terms of message sharing, on average.

The results in Figure 5.7 show great diversity in social curiosity stimulation profiles across different users and, for many users, very dynamic profiles. To further illustrate this point, we show in Figure 5.8 the time series of both social curiosity metrics – $maxDirInf$ e $maxIndInf$ – for six selected (user, group) pairs. These curves illustrate how the social curiosity of particular users in specific groups evolve over time, as captured by both metrics. To facilitate visualization, the x -axis of Figures 5.8a–5.8f represent the normalized *lifespan* of the particular user in the given group, that is, the time interval between the first and last message shared by this user in the group, normalized to fall between 0 and 1.

Overall, we can see some important fluctuations over time, notably in terms of the maximum *direct* influence¹⁶. Such fluctuations follow the natural dynamics of user participation in the group. Figures 5.8a and 5.8b show the time series of the social curiosity of the same user (identified as user 0) in two different groups she participates in (identified as groups 1 and 2). We note very different patterns, temporally and on average. Indeed, the dynamics captured by these two sets of curves fall into two very different profiles, identified by UBMGs 2 and 1, respectively. This observation hints at the validity of one of our key assumptions, that is, that user curiosity stimulation varies depending on the group. Figure 5.8c shows the time series of another user in a different

¹⁶As already shown in Figure 5.3b, the values of $maxDirInf$ tend to be much larger than those of $maxIndInf$.




Figure 5.9: Diversity of social stimulation of selected users in different groups: the same user in different groups is represented by the same color.

group. As illustrated in the figure, this user’s social curiosity fluctuates greatly over time, with alternating behavior of greater and lower stimulation.

Figures 5.8d–5.8f shows the time series of three different users, members of the same group. For that reason, the x -axes of these three graphs represent the normalized lifespan of the given *group* (i.e., identified as group 4), that is the time interval between the first and the last message shared by any user in the group, normalized to fall between 0 and 1. This group was selected as one among those with the largest number of distinct users during the monitoring period. Once again, we see that user social curiosity fluctuates greatly over time (notably user 3, shown in Figure 5.8e). Most importantly, we see very different patterns across the three users, even though they are members of the same group. This observation hints at the idea that the impact of the group on the curiosity of individual members may be different for distinct users. We further elaborate on this point in the following section, as we present our key results on group-level social curiosity stimulation.

5.5.2.3 Social Curiosity at the Group Level

To analyze social curiosity at the group level, we focus on two aspects, notably: (1) the role of the group on the curiosity stimulation of its members and (2) the overall social curiosity of each group. For the latter, we use the group-level social curiosity metrics

Table 5.6: Test of statistical difference of social stimulation metrics of the same user on different groups (users 0–7 refer to selected users in Figure 5.3b).

User	Total # Msgs	# Groups	# Pairwise Comparisons	Percentage of Pairs of Groups with Different Metrics	
				Max. Direct Influence	Max. Indirect Influence
0	4,361	17	136	71 (52.21%)	70 (51.47%)
1	1,210	13	78	15 (19.23%)	11 (14.10%)
2	1,736	15	105	75 (71.43%)	75 (71.43%)
3	2,459	12	66	13 (19.70%)	25 (37.88%)
4	1,168	15	105	32 (30.48%)	31 (29.52%)
5	976	6	15	4 (26.67%)	5 (33.33%)
6	803	7	21	5 (23.81%)	11 (52.38%)
7	558	11	55	14 (25.45%)	3 (5.45%)

defined in Section 5.4.

To investigate the role of the group on the curiosity of its members, we selected the 8 users who shared the largest number of messages in multiple groups, and computed the *average* values of both social influence metrics for each (user, group) pair. Figure 5.9 shows the results as a scatter plot where each circle is a (user, group) pair. The same user (in different groups) is shown in the same color and the circle’s diameter represents the number of messages shared by the user (in the given group). Visually, there are great distinctions on both metrics (in terms of average values) for the same user depending on the group she is participating in.

These distinctions were confirmed by pairwise statistical tests of the differences for each of the two metrics. Specifically, for each user u , we used a non-parametric statistical test capable of identifying the significant differences between tests, namely the Kruskal-Wallis Honestly Significant Difference (HSD) test [115] (with $\alpha=0.05$), to test the difference between the maximum direct influence for pairs of groups in which u participates. We did the same for the maximum indirect influence. The results are reported in Table 5.6, which shows, for each selected user u , the total number of messages shared by u , the number of groups u participates in, the total number of pairs tested and the number (and percentage) of pairs for which there is a statistical difference in each metric (two rightmost columns). As shown, the differences are statistically significant for a large fraction of the cases for several users. Thus indeed the social stimulation of the same user may be quite distinct depending on the group she participating in, as already suggested by the results in Figures 5.8a and 5.8b. Indeed, we note that user 0 in Table 5.6 refers to the same user 0 shown in those two figures.

As a final analysis, we use the group-level social curiosity stimulation metric, defined in Equation 5.12, to characterize the groups in our dataset. Recall that this metric captures a reduction in uncertainty associated with curious users, defined in Equation 5.11, due to the knowledge of the social influence from other users (stimulating users) in the group. Thus, to facilitate analysis, Figure 5.10 shows a scatter plot with the values




Figure 5.10: Social curiosity stimulation across groups: group mutual information versus entropy of curious user, measured at the the time of the last message in the group (groups of user 0 in blue).

of these two metrics – group mutual information ($groupMutInf$) and entropy of curious user – computed for each of the 335 groups at the time of the last message sharing in each group.

The figure shows great diversity in social curiosity also at the group level. Many groups have a value of $groupMutInf$ (x -axis) very close to the entropy of curious user $H_{t,g}^{\leftarrow}(U)$ (y -axis), suggesting that a large fraction of the total uncertainty associated with the curious users in those groups can be captured by social influence. In other words, social influence is a key component of curiosity stimulation in the group. This happens for groups with distinct profiles in terms of number of users and level of activity of these users, which ultimately lead to groups with distinct values of entropy of curious user. On the other hand, a number of groups have values of $groupMutInf$ much smaller than the total uncertainty associated with the curious users. In those cases, social influence plays a less important role on curiosity stimulation at the group level.

For illustration purposes, Figure 5.10 also highlights (in blue) the groups in which one selected user – user 0 in Figure 5.9 – participates. As shown, the diversity in user 0’s social stimulation shown in Figure 5.9 follows the diversity present also in his groups. We measured the correlation between group-level social curiosity and individual-level social curiosity (as captured by both metrics considered) for the selected 8 users in a few selected time windows but found not clear pattern: some cases of positive correlation, several cases of negative correlation and some cases of no-correlation at all. Thus, even though user curiosity stimulation does indeed vary depending on the group, we found no consistent indication that a more socially stimulated group always translates into the same behavior on all its members. Indeed, as shown in Figures 5.8d–5.8f, different members of the same group may exhibit quite distinct patterns of social curiosity. More broadly, it might be




Figure 5.11: Group-Level Social Curiosity Stimulation: CDFs of (a) group lifespan and (b) the impact of social curiosity, measured when group reaches 10th, 50th and 90th percentiles of user participation; (c) Time series of the impact of social influence of 3 example groups; (d)–(f) Time series of group-level metrics for the same 3 groups.

the case that different users dominate the overall group-level curiosity stimulation pattern at different time windows.

Recall that Figure 5.10 shows results for a specific time window, that is, the time of the last message sharing in each group. We finish our discussion by showing the dynamics of the group-level social curiosity metric over time. Naturally, one key component of social curiosity at the group level is user participation in the group. Thus, we analyze social curiosity in each group in different points in time defined based on user participation. Specifically, for each group g , we looked at the total number of *distinct users* who shared some message in g since the first message until the last message shared during the monitoring period, and considered the time when g reached the i^{th} percentile of user participation¹⁷. We then looked at the lifespan of each group, and measured the time when each group reached the i^{th} percentile of user participation, as a fraction of its total lifespan.

Figure 5.11a shows the Cumulative Distribution Functions (CDFs) of the fractions of group lifespans associated with the 10th, 50th and 90th percentiles of user participation in each group. Whereas some groups reach all three percentiles quickly, others take

¹⁷Obviously, group activity prior to the monitoring period is disregarded. Thus, the analysis reflects group activity during this period only.

much longer to attract greater diversity in user participation, reaching the 90th and even the 50th percentiles very later on during monitoring. For instance, roughly 40% of the groups take more than 50% of their lifespans to reach the 90th percentile of user participation. In light of this dynamics and heterogeneity, Figure 5.11b shows the CDFs of the *impact of social curiosity* which is the ratio between the reduction in entropy of curious user due to social influence and the entropy of curious users, is defined as the ratio $groupMutInf(t, g) / H_{t, g}^-(U)$, measured at the same points in time (same three percentiles). Recall that the greater the impact of social curiosity of curious user the more important the role of social influence on curiosity stimulation at the group level. Once again, there is great diversity across groups, here shown over time. Roughly 30% of the groups have impact of curiosity at least 50% at the time they reach the 10th percentile of user participation. In contrast, about 70% of groups at the 50th percentile have the same (or less) impact of social curiosity. In addition, only 10% of groups at 90th percentile have between 50% and 100% of impact of social curiosity.

Moreover, groups may experience quite dynamic social curiosity stimulation, evolving over time according to different patterns for distinct groups, as an aggregation of the individual patterns of the current group members (which, as discussed in the previous section, is also quite dynamic and heterogeneous). This is illustrated in Figure 5.11c, which shows the time series of the impact of social curiosity for three selected groups. Such patterns illustrate the very dynamic role of social influence as a driving force behind curiosity stimulation in the group: in some cases, social influence starts as an important component of group-level curiosity and continuously decreases over time (e.g., group 6 in the figure); while in other cases, the importance of social influence fluctuates over time (e.g., group 5), possibly reaching rough stability (e.g., group 7).

In order to explain such varied patterns, we plot in Figures 5.11d–5.11f the three metrics associated with group social curiosity, namely the $groupMutInf$ (represented by the red line), $groupEntropy$ (represented by the green line) and entropy of curious user, or $H_{t, g}^-(U)$ (blue line) for the same three selected groups. For the groups 5 and 7 shown in Figures 5.11d and 5.11f, social influence plays an important role on the overall group curiosity over large periods of time, as the $groupMutInf$ curves represents a large fraction of the total entropy of curious user (over 60%, on average). For both groups, $groupMutInf$ (red line) decreases mostly when the $groupEntropy$ (green line) increases, that is, when the probabilities of the same *curious user* conditioned on the same *stimulating user* decrease. In that case, the chance of the same users co-occurring in the same window decreases. Thus, social influence loses relevance as a curiosity driver for the group. When this happens at the same time as an overall drop in entropy of curious user, possibly due to an increase in the number of distinct users sharing messages, the role of social influence on curiosity at the group level becomes even less important (e.g., group 5 at around 50% of its lifespan). In contrast, Figure 5.11e shows a higher $groupEntropy$

Table 5.7: Three example messages shared by the most active user in group “Science, Religion and Politics”.

Msg Id	Curiosity Profile	Max. Direct Influence	Max. Indirect Influence	Window of Interaction	
				# messages	# users
1	Dependent	6.78	0.03	28	6
2	Indirect	0.50	0.10	61	7
3	Independent	1.50	0.08	90	8


(green line) which dominates the total entropy of curious user for group 6 throughout its lifespan. Clearly, the curiosity of this group, in general, is very weakly impacted by social influence, as shown in Figure 5.11c.

5.5.2.4 Some Illustrative Examples

In this section we discuss a few examples to further illustrate the use of our metrics. We start by looking at the participation of a particular user, referred to as user u , in a group entitled “Science, Religion and Politics” (translated to English), which is one of the groups with the largest number of distinct members (353 users during the complete monitoring period) and messages shared (48,389). User u is the most active user in this group, having shared a total of 3,787 messages. Also, u ’s curiosity stimulation follows profile U_0 which, as shown in Table 5.5, experiences all three states of message-level curiosity stimulation. Thus, we selected three messages shared by u , referred to by identifiers 1, 2 and 3, associated with each profile – dependent, indirect and independent, respectively.

Table 5.7 presents information associated with each selected message, notably its curiosity stimulation profile, the values of two social curiosity metrics (maximum direct and maximum indirect influences) as well as statistics of the window of interaction defined when the given message was shared. The latter includes the number of distinct users (i.e., stimulating users of influence) and number of messages shared during the period.

We note that the number of potential sources of social influence on u when she shared message 1, i.e., the number of stimulating users of influence, is the smallest among the three messages (6). The same holds for the total number of messages shared during the window of interaction (28). Yet, this message has the largest value of maximum direct influence (6.78), much larger than the other messages, and the smallest value of maximum indirect influence (0.03). Thus, despite in smaller number, the past behavior of these stimulating users and of the curious user u offer strong evidence of *direct* social influence driving u ’s content sharing decisions. This message is characterized by the



(a) Group: “We are all HADDAD!”.

(b) Group: “Lula the father of prisoners”.

Figure 5.12: Graph representation of direct social influence among members of two selected groups (node diameters are proportional to out-degrees, red/blue edges refer to strong/weak social influence, thus strong/weak social curiosity stimulation).

dependent profile.

In contrast, when sharing message 2, user u 's curiosity is not strongly stimulated by those sharing content during the window of interaction, at least not with respect to the stimulus triggered by direct influence estimated by prior interactions. In other words, the maximum direct influence is very small (only 0.5). Yet, compared to message 1, we do find more evidence of indirect influence from the stimulating users stimulating u 's curiosity at this time: the maximum indirect influence has the largest value of the three messages, and the message is characterized by the *indirect* profile. Indeed, some of these stimulating users are among the top 10 most active users in the group. Also, we tracked these stimulating users over successive windows, observing an increase in the *direct* influence from some of them on u , which suggests that, as mentioned in Section 5.4.2, sources of indirect influence may become stronger sources of direct influence as time progresses.

Finally, when sharing message 3, user u 's curiosity is not strongly stimulated by neither direct nor indirect influence, despite the presence of large numbers of messages and users in the corresponding window of interaction (the largest ones out of the three messages). The message is characterized by an *independent* profile. Clearly the more intense user activity during the window does not translate into stronger curiosity stimulation on u : both social curiosity metrics have intermediate values.

To illustrate social curiosity at the group level, we build a graph representation of social influence based on the contingency table, which contains historical patterns (see definition in Section 5.4.2). In this graph representation, each node represents a group member and a direct edge from user s to user u is added to represent that s

(stimulating user) has influence on u (curious user). Edge weights are given by the conditioned probability of a curious user given an stimulating user, i.e., $P_{t,g}^+(U=u | S=s)$, computed based on values in the contingency table.

Figure 5.12 shows the graphs built for two specific groups considering the complete history of user interactions in each group, that is, by the time the last message is shared in the group. In each graph, node diameters are proportional to out-degrees and edge colors denote the strength of the conditioned probabilities used as weights (and thus the strength of the social influence computed from such probabilities). Specifically, red edges are those whose weights are above 0.1, thus representing stronger social curiosity stimulation, and blue edges have weights below 0.1, suggesting weaker social curiosity stimulation.

Figure 5.12a represents a group entitled “We are all HADDAD!”, which had 147 distinct members who shared 3,012 messages during the monitoring period. Fernando Haddad was a left-wing presidential candidate in the 2018 general elections in Brazil. Approximately 95% of the edges in the graph are blue, suggesting that social influence does not play an important role in the group. Indeed, the impact of social curiosity is below 6%, i.e., the overall reduction of entropy of curious user due to social influence. This happens despite the great number of edges connecting users, that is, there is some evidence of social influence, but in the vast majority of cases this evidence is very weak. Indeed, the authors are aware that most members of this group often shared news about the candidate’s campaign, which did not generate much response or engagement from most members. The occasional user interactions resulting from such messages offer only weak evidence of influence (if any at all). However, there are a few cases of more responsive and intense discussions on specific political themes, often with the participation of particular users. These cases show up in the few red edges in the top part of the graph. Note that many of them are linked to a particular node, possibly a user who is frequently involved in the discussions.

Figure 5.12b, in turn, shows the graph representation for the group entitled “Lula the father of prisoners”, which had 221 distinct members and 935 messages shared during the monitoring period. Lula is a former president of Brazil and a charismatic representative of the left wing who was arrested during the period of the 2018 general elections. Unlike in Figure 5.12a, this graph is very balanced in terms of the strength of social influence among users: 51% (49%) of the edges are blue (red). Overall, social influence is a much stronger component of curiosity stimulation, with a impact of social curiosity of 62% (i.e., reduction of entropy). This group is mostly driven by very passionate pro-right wing discussions, thus the group’s title is a sort sarcasm or irony. As reflected in the graph, such discussions, often involving the same participants, do offer strong evidence of curiosity stimulation as a driver behind content sharing. Note also that, compared to the graph in Figure 5.12a, the graph in Figure 5.12b is much more sparse, suggesting that such strong curiosity stimulation occurs between specific pairs of users for whom the bond

Table 5.8: Validation of users.

Description	Users (Percentage)
1 Beta	6,571 (79.11%)
2 Beta mixture	1,098 (13.22%)
3 Beta mixture	239 (2.88%)
No valid users	398 (4.79%)
Total of users	8,306 (100.00%)

created by prior interactions offer greater evidence of social influence.

5.6 Curiosity Models

In this section, we employ the methodology applied to the LastFM dataset of the first case study in Section 4.4 to estimate the Wundt’s curve, the curiosity model that represents the curiosity profile of a user.

We use the Beta mixture distribution defined in the Equation 4.9 with the shape parameters $\alpha > 1$ and $\beta > 1$ to capture the complex behavior of the users from the WhatsApp case. For this purpose, we also consider $M \in \{1, 2, 3\}$, where M is the number of mixture components for the Beta mixture distribution, and map our final stimulus values obtained by our metrics into the range $[0, 1]$, which is the domain of the Beta distribution.

For each user, we calculate the simple mean of the metrics for each stimulus value of user. So, we consider the social curiosity variables (i.e, *Max Direct Influence* and *Max Indirect Influence*) and the traditional collative variables (i.e., *User Novelty*, *Category Uncertainty* and *User Uncertainty*) extracted in Section 5.5.1 after complementarity analysis of all metrics. We then mapped these mean values to the range $[0, 1]$ using the min-max normalization to compose the final stimulus value.

5.6.1 Classes of Wundt’s Curve

We fit the Beta mixture distribution to the curiosity model by considering three values of M for each user in the WhatsApp dataset and using the KS test to validate this fit. Furthermore, since we perform multiple hypothesis testing, we use the BH method with a significance level of $\alpha^+ = 0.01$ to adjust the p -values for each of the multiple




Figure 5.13: Mixed Wundt's curve: examples of 4 users.

comparisons of the KS test. In the same way, we consider the smallest value M for the Beta mixture distribution validated by the KS test when fitting Wundt's curve.

Table 5.8 shows the results of the fit for the curiosity model and it can be seen that 79.11% of users are accurately modeled by a single Beta distribution, while 13.22% of users are modeled by $M = 2$ Beta mixture and 2.88% by $M = 3$ Beta mixture. Only almost 4.8% of users were not validated with any value of M and are therefore discarded¹⁸.

Finally, Figure 5.13 shows an example of the fitting of the mixed Wundt's curve to four users and thus the complexity of their user behavior due to the multi-modality of the curve. Figures 5.13a and 5.13b show the users whose curiosity model was estimated with a 2-Beta mixture, while Figures 5.13c and 5.13d show the users with a 3-Beta mixture.

5.6.2 Curiosity and Stimulation Patterns

Table 5.9 shows the fraction of the stimulus profile within each Beta mixture modeled by the curiosity model. Each column of the table sums to 100% and the fraction for each Beta mixture is calculated over all stimulus values for each user with single Beta,

¹⁸It is interesting to note that the use of social curiosity and the traditional variables shows the contribution of social influence as complementary variables to the formation of the final stimulus and to the estimation of Wundt's curve. In contrast, using the social curiosity variable in isolation when composing the final stimulus leads to an increase in the number of invalid users from 4.8% to 33%.




Figure 5.14: CDF of fraction of social curiosity stimulation per profile per user.

2-Beta and 3-Beta mixture. It can be seen that the profile *Independent* has the largest fraction regardless of the type Beta mixture, namely than 70%. This corresponds to the overall distribution for this profile according to Table 5.4 in Section 5.5. The stimulus profile of *Indirect* and *Dependent* also follows the overall distribution for their respective profiles, with the exception that the fraction of *Dependent* in the 2-Beta mixture is greater than that of *Indirect*, namely 15.21% versus 13.57%. Although the fractions for *Dependent* and *Indirect* are approximately equal.

If we consider only these fractions of the stimulus profile per Beta mixture models of the Table 5.9, it is not possible to identify a clear pattern for the differences between the profiles and the user's mixed Wundt's curve. So, to better understand the relationship between the mixture models and the stimulus profiles when looking at individual users, we calculate the stimulus value per stimulus profile normalized by user as follows to show their distributions as follows.

Table 5.9: Statistics for profile of stimulation per Beta (mixture) distribution.

Stimulus Profile	1 Beta	2 Beta mix	3 Beta mix
<i>Independent</i>	73.19%	71.22%	72.56%
<i>Indirect</i>	15.35%	13.57%	14.22%
<i>Dependent</i>	11.46%	15.21%	13.22%
Total:	100.00%	100.00%	100.00%

Figure 5.14 shows the CDF where each curve represents an stimulus profile (i.e., *Independent*, *Dependent* and *Indirect*) and each point in the curve is a user. Figures 5.14a, 5.14a and 5.14a show the single Beta, 2-Beta and 3-Beta mixture models respectively. Regardless of the figure, the curve of the *Independent* profile (blue solid line) show the same behavior. In contrast, in Figure 5.14a the *Dependent* stimulus profile has a share of 75% of users up to 25%, while in Figure 5.14b this same share of users is up to 35% and in Figure 5.14c up to 33%. This shows that the *Dependent* profile increases when the users are modeled by a 2-Beta and 3-Beta mixture.

The *Indirect* profile (dashed black line) in Figures 5.14a and 5.14b shows a wider distribution of fractions than the *Dependent* profile (dashed-dotted red line) for the single Beta and the 2-Beta mixture. In contrast, *Indirect* has a smaller distribution than *Dependent*, i.e., the dashed black line lies above the dashed-dotted red line in Figure 5.14c




Figure 5.15: Profile of social curiosity stimulation: examples of four users.

for the 3-Beta mixture. In summary, this means that while the fraction of the *Dependent* profile increases for the Beta mixture models, the fraction of the *Indirect* profile increases for the 3-Beta mixture, where the users exhibit more complex behavior for their curiosity profile (multi-modal Beta mixture distribution).

In addition, Figure 5.15 shows the four users of the mixed Wundt's curve example from Figure 5.13 with histograms separated by stimulus profile. The stimulus values of the four users are concentrated on the profile *Independent* (purple histogram), which is present in every figure of this figure. Furthermore, the profile *Independent* is balanced with respect to the profile *Indirect* in all mixed Wundt's curves. While the profile *Indirect* is only present in Figure 5.15d (yellow histogram), one can easily notice that the profile *Dependent* (green histogram) is more pronounced than the profile *Indirect*, which corresponds to the same behavior of the CDF of Figure 5.14c, as mentioned above.

What is the most interesting is how in the Figures 5.15a the different stimulus profiles, such as *Independent* and *Dependent*, form different histograms that do not mix, i.e., the User 1 has two very clear stimulation patterns. This rises an important question: *How does the user's curiosity profile, modeled by a mixture of Betas, evolve over time?* This question is addressed in the next subsection.




Figure 5.16: Example of user with 2-Beta mixture: Wundt's curve and stimulation profile.

5.6.3 Temporal Analysis

To complete the analysis of curiosity models in this section, a temporal analysis of the mixed Wundt's curve for 2-Beta and 3-Beta mixtures is performed. To this end, we investigate whether the user's curiosity profiles are always estimated by the Beta mixture even for short time periods. The methodology used here is the same as in Section 4.4.3 of the case study for the LastFM dataset. In this way, a sample of 25 users for the 2-Beta mixture and a sample of 25 users for the 3-Beta mixture are drawn at random from the dataset list of users with at least 1,000 stimulus values. We emphasize that the 25 selected users of the 2-Beta mixture represent 17.2% of the users modeled by the 2-Beta mixture distribution (i.e. 1,098 users – Table 5.8) and have between 1,000 and 24,200 listening events. The 25 selected users of the 3-Beta mixture represent 11.72% of the users modeled by the 3-Beta mixture distribution (i.e. 239 – Table 5.8) and have between 1,000 and 14,700 listening events. A total of 50 users were therefore selected.

The user's activity period was divided into short time periods, as shown in Figure 4.11 of Section 4.4.3, and for this analysis we specify the short time periods as one week. Also, recall that Wundt's curve is estimated for each short time period (i.e., one week) of the user. However, each short time period should contain at least 30 stimulus events




Figure 5.17: Example of user with 3-Beta mixture: Wundt's curve and stimulation profile.

for the curiosity model to be estimated. The metrics used to evaluate the experiments in this analysis are $avgPrecision_{\beta}$ and $avgPrecision_{AP}$ which were defined in the Equations 4.10 and 4.11.

Table 5.10 shows the results of the experiments for the temporal analysis. In this table, the metrics are presented with a confidence interval of 95%. It can be seen that for the metric *average precision of estimated Beta mixture distributions* ($avgPrecision_{\beta}$) more than 96% of users have a more complex curiosity profile, i.e., more than a single Beta distribution is required to estimate Wundt's curve for each period of a week. However, for most users (i.e., more than 54%), the metric *average precision of the stimulus profile* ($avgPrecision_{AP}$) shows that each period of a week contains a different stimulus profile.

Table 5.10: Results of experiments for temporal analysis with C.I. of 95%.

Metrics	2 Beta mixture	3 Beta mixture
$avgPrecision_{\beta}$	0.968 ± 0.042	0.986 ± 0.014
$avgPrecision_{AP}$	0.763 ± 0.072	0.649 ± 0.101

To illustrate our analysis, Figures 5.16 and 5.17 show a temporal analysis for two users, the first modeled by a 2-Beta mixture and the second modeled by a 3-Beta mixture distribution. These figures are arranged so that one line shows Wundt's curve for each one-week period and next line the histograms of stimulation profile is shown for each

Wundt's curve.

Figure 5.16 shows a user with 9 periods of a week and $avgPrecision_\beta = 1.0$, which means that for 100% of his periods were estimated by a mixture of the Beta distribution. It has all periods with at least two stimulus profiles, i.e., $avgPrecision_{AP} = 1.0$. In addition, 7 out of 9 weeks have all the stimulus profiles (i.e., *Independent*, *Dependent* and *Indirect*). The user in Figure 5.17 also presents 9 periods of a week and has $avgPrecision_\beta = 1.0$ and $avgPrecision_{AP} = 1.0$. In contrast to the first user, this user only has the profiles *Independent* and *Dependent* and shows multi-modal Beta mixture in all weeks, i.e., he has a dynamic behavior with regard to the curiosity stimulation.

5.7 Validation of Curiosity Model

In this section, we use null models to validate our metrics for the social curiosity model. Then, we use them for comparison purposes to verify that the social curiosity metrics actually capture user's social curiosity behavior on the group communication platform. To this end, we use the additional assumptions for social influence in WhatsApp defined in Chapter 3 as a basis to generate a set of hypotheses to validate our social curiosity model.

We recall these additional assumptions: (1) that a user's curiosity may be triggered differently depending on the people with whom he or she interacts and the ongoing discussions between them; (2) that the curiosity of the user who shares certain content can be stimulated by other users who shared content in the same group during the window of the interaction through social influence; and finally (3) the extent to which the user's curiosity is stimulated by other users can be estimated from the historical patterns of their action events.

The hypotheses for testing the null models are as follows:

Null hypothesis 5.1 (H5.1): *Social curiosity is influenced by interpersonal interest.*

Null hypothesis 5.2 (H5.2): *Social curiosity is influenced by temporal constraints such as recency and temporal locality – user's tendency to interact with the same user repeatedly over time.*

Null hypothesis 5.3 (H5.3): *Social curiosity is influenced by the frequency and the total number of sharing messages.*




Figure 5.18: Design of Null Model 0 and 1.

From this set of hypotheses to be tested, each of the additional assumptions of social curiosity can be related as follows: (i) **H5.1** refers to the first additional assumption, whose focus is on *interpersonal interest*; (ii) the second additional assumption, in turn, refers to **H5.2**, which deals with the *window of interaction*, that is, temporal constraints; (iii) **H5.3**, finally, refers to the third additional assumption, which concerns the *historical pattern* of the user's action events.

To avoid long execution times in the validation experiments with the null models, we selected the first 100 users for whom the number of shared messages was above the 90th percentile of the number of messages in the group, i.e., users with more than 1,000 shared messages. In addition, these users should be members of groups with less than 30 thousand shared messages and the ratio between *number of messages of the user* and *total number of shared messages of the group* should be more than 10% and less than 60% (to avoid that one selected user dominates the conversations in the group).

The Figure 5.18 illustrates the design of the *Null Model 0* and *Null Model 1*. On the right side of this figure, the target user (indicated by the blue shaded area) is retained, while the users in the window of interaction are randomly selected from the set of influencer users with whom the target user has interacted in the past. This random selection is according to: (a) the distribution for the target user's interaction preference for *Null Model 0*; or (b) the uniform distribution for *Null Model 1*. The first distribution is calculated using the frequency of influencer users with the target user in the historical patterns, and the second considers the same probability for each influencer user. Furthermore, the timestamps for both random samples of the null models are identical to those of the real samples. This is illustrated in Figure 5.18 by the dashed vertical lines connecting the timestamps of the real sample (horizontal black line at the top of the figure) and the random sample generated by the null model (horizontal red line at the bottom of the figure). Thus, for each shared message by the target user and their respective window of interaction, influencer users are randomly selected according to its respective null model. This leads to a change in the historical patterns for the group conversations




Figure 5.19: Design of Null Model 2.

and especially for the target user. Finally, by randomly selecting users in the target user's window of interaction, *Null Model 0* aims to break the *temporal constraints* and thus tests **H5.2**. *Null Model 1*, in turn, aims at breaking the *interpersonal interests* and *temporal constraints* and thus tests **H5.1** and **H5.2**.

The Figure 5.19 shows the design of the *Null Model 2*, which aims to test **H5.3**, that is, to breaks the *frequency of influencer users* and the *total number of messages shared* into historical patterns. This is an extension of the *Null Model 1*, but with one change: the total number of users in the window interaction is randomly generated according to the total number of users preferred by target user. For this purpose, the distribution of the target user's preferences over the total number of users within the window of interaction is calculated based on the frequency of the total number of users considered in each window interaction for the target user. In this way, the timestamps within the window of interaction are generated evenly according to the total number of previously selected users. Then, the influencer users are uniformly selected according to the set of influencer users that the target user has interacted with in the historical patterns. The Figure 5.19 shows that the total number of users in the random sample timeline (horizontal red line at the bottom of the figure) is larger than the total number of users in the real sample timeline (horizontal black line at the top of the figure). So, there are 12 influencer users in the random sample versus 9 influencer users in the real sample. Finally, in this figure the timestamps for real sample and the random sample are different, except for the timestamp of target user, which is the same (i.e., the timestamp t_{10} is equal to t_{13} on the right side of Figure 5.19).

To make a comparison between the real stimulus values measured by the users (real sample) and the stimulus values randomly generated by each of the null models (random samples), we calculate the difference between the real stimulus value and the random stimulus values. Then, positive values of these differences indicate that the real stimulus values are larger than the random stimulus values of the null model. In contrast, negative values of the differences indicate that the random stimulus values of the null




Figure 5.20: CDFs of social metrics for Null Models, *Maximum Direct Influence* and *Maximum Indirect Influence*, respectively. Each one of CDFs denotes the identified message-level social curiosity stimulation profiles: (a) *Independent* depicted by blue line; (b) *Indirect* represented by red line; at end, (c) *Dependent* illustrated by dashed black line.

model are larger than the real stimulus values. The set of CDFs in Figures 5.20 shows these differences to the *Maximum Direct Influence* and the *Maximum Indirect Influence* at the social curiosity stimulation profile level identified in Section 5.5.2 for all 100 users selected to know (a) *Independent*; (b) *Indirect*; and (c) *Dependent*.

The Figures 5.20a–5.20c show the difference distributions of the *Maximum Direct Influence* for each of the three null models – *Null Model 0*, *Null Model 1*, and *Null Model 2*, respectively. We note that the distribution of this metric does not change for the *Independent* profile (blue line) and remains stable for any null model. This is because in this profile all metrics have smaller values than in other profiles. However, this is not the case for the other profiles. For the *Indirect* profile (red line), for example, a slightly more than 60% of the values in the distribution for all null models are negative, which means that the randomly generated stimulus values by null models are larger than the real stimulus values. Whereas the *Dependent* profile (dashed black line) has more than 80% positive values, i.e., the random stimulus values are larger than the real stimulus values. However, if we look at the results of the distribution of differences from *Null Model 0* to *Null Model 2*, we can see that these distributions show an smooth increase, which means that the differences for both profiles, the *Indirect* and the *Dependent*, are increasing.

Like the previous figures, Figures 5.20d–5.20f show that the distribution of dif-




Figure 5.21: CDFs of stimulus values of *Maximum Direct Influence* and *Maximum Indirect Influence* for *Dependent Profile* and *Indirect profile*: (a) Real stimulus values are depicted in *green line*; (b) random stimulus values of Null Model 0 denoted by *blue line*; (c) random stimulus values of Null Model 1 in *red line*; at last, (d) random stimulus values of Null Model 2 in *dashed black line*.

ferences for *Maximum Indirect Influence* and the differences for the *Independent profile* (blue line) are constant and do not change. Nevertheless, the *Indirect profile* (red lines) in these figures shows the most of positive differences, more than 60% for the *Null Model 0* results and almost 93% and 98% for *Null Model 1* and *Null Model 2*, respectively. For *Dependent profile* (dashed black line), they show that almost 60% of the differences are negative values (in contrast to the *Maximum Direct Influence* distributions). As in the previous figures, the distributions show an increasing trends for the distributions of *Null Model 0*, *Null Model 1*, and up to *Null Model 2*, i.e., the more the null models change, the further the randomly generated stimulus values are from the real values.

It is important to note that each of the proposed null models is further from the user's preference and thus closer to a random pattern. This means that the difference from the real stimulus values gradually becomes larger. Yet, we emphasize that for the *Dependent profile* the *Maximum Direct Influence* is predominant compared to the other profiles, similarly, the *Maximum Indirect Influence* is dominant for the *Indirect profile*. Thus, depending on the stimulation profile of social curiosity, the differences in the metrics between the real stimulus values and the random stimulus values of the null model are more or less pronounced.

The Figure 5.21a shows the *Maximum Direct Influence*, i.e., the metric that most

influences the *Dependent* profile, and Figure 5.21d shows the *Maximum Indirect Influence*, i.e., the metric with the greatest influence on the *Indirect* profile. In these figures, we can see that the curve for the distribution of differences from *Null Model 2* has the lowest stimulus values, followed by *Null Model 1* and *Null Model 0*, respectively, and this fact occurs in both profiles because they are the dominant metric in their respective profiles.

In contrast, Figures 5.21b and 5.21c show that the metrics that are not dominant in the respective profiles exhibit different behavior, i.e., the real stimulus values (green dashed line) are lower than the random stimulus values generated by null models. So, in decreasing order, the distributions with higher values were generated by *Null Model 0* (blue line), *Null Model 1* (red line) and *Null Model 2* (black dashed line), respectively, and it is the opposite behavior when compared with the dominant metric values. This means that the social influence generated by the null models is somewhat higher than the measured real stimulus values. In other words, for the profile *Dependent*, the social strength of the most social indirect *stimulating user* in the group was increased by the way the random sample of each null model was constructed, which is related to the historical pattern. Also, for the *Indirect* profile, the social strength of the direct *stimulating user* for the *curious user* was increased by the way the random sample of the null model was constructed, which is related to the historical pattern.

Overall, our null hypothesis study shows that our stimulation values for the majority of users and especially for the *Dependent* and *Indirect* profiles are not explained by chance (null models). Thus, our stimulation values measure a correlation that is present in the WhatsApp dataset.

5.8 Limitations of our Study and their Implications

In the process of modeling such a highly subjective and complex concept such as social curiosity, we made a number of assumptions and simplifications so as to be able to derive metrics that are reasonably simple while still useful and informative. In the following, we discuss some of these limitations and their implications.

5.8.1 Lack of Content Representation

Message content is most probably an important factor influencing how one's curiosity is stimulated. Specifically, the contents of messages posted by others may indeed be a component of how *social* curiosity affects one's behavior. However, we chose not to exploit it in the derivation of our *social* curiosity metrics, focusing only on the primary factors, which, we argue, are related to who is sharing content in the group, i.e., users and their (inferred) social links. There are several reasons to do so:

1. Access to message content is not always available. Thus, by not using message content we assure that our metrics of social curiosity can be used even when content is not accessible (as is the case of our dataset).
2. As mentioned, by focusing only on the users sharing content, we aim at addressing a novel aspect of curiosity stimulation that is not captured by the other (traditional) collative variables. Thus, we complement these variables by proposing metrics related to a novel aspect.
3. By not including content in the social curiosity metrics, we are able to assess the extent to which the most primary component of social influence (i.e., people themselves) can drive one's curiosity.

We note however that, even though the content factor is not explicitly included in the proposed metrics, these metrics do exploit prior user interactions, which in turn, may have been influenced by the content shared by those users at those prior moments in time. Thus any content related effect that might impact social influence is being indirectly captured when we make use of such user interactions.

Having argued for not exploiting content to estimate social curiosity, we note that we did use message categories as a proxy for content properties when deriving metrics related to the traditional collative variables (see Section 5.4.3, acknowledging its importance as a component of curiosity stimulation in general – not only from a social perspective). Obviously, categories, as those used by us, offer only a very coarse approximation of content properties and their use as such is another a limitation of the study, as we further elaborate next.

5.8.2 Different Dimensions of Online Information Dissemination at the Level of User Behavior

We are aware that there are several dimensions that lead an individual (i.e., a user) to share an information content in the online information dissemination process that go beyond social curiosity and that are not addressed in this work, namely: complex contagion, polarization, opinions and beliefs/creeds.

In a complex contagion, the success information dissemination depends upon (social affirmation) interaction with multiple sources. That is, multiple sources of exposures are required before an individual adopts a change of behavior [26, 49]. For example, most tweets disseminate through complex contagions in retweet networks, moreover, political hashtags behave like complex contagions, while news-based hashtags spread like simple contagions [26]. In addition, there are authors who emphasize that heterogeneity of adoption can be widespread in many information dissemination process due to individual diversity of persistence, creed, and preexisting information [103]. Thus, the question of how social curiosity affects complex contagion in different types of online social platforms is a complementary study that we chose do not tackle in this chapter.

It is known that polarization exists mainly in mass media and social, interaction and collaboration networks (e.g., online social networks) and it is characterized by an increasing agreement intragroup (i.e., members with similar beliefs) and a deepening disagreement intergroup (i.e., members with contrary beliefs) [56]. Yet, according to Interian *et al.* there are two main factors that shape people's opinions: confirmation bias and social influence [56]. Social influence is therefore also the process by which individuals adapt their opinions, revise their beliefs or change their behavior as a result of social interactions with other people [107]. As we have focused on social influence as the main driver of social curiosity in this chapter, we emphasize that the polarization, the formation of opinions and beliefs can directly influence social curiosity and that we cannot identify its role in stimulating curiosity due to the lack of representative content (previous Section 5.8.1) and the lack of metadata to verify whether the user actually read the messages (see details in Section 5.8.5). Although there are several dimensions in the process of online information dissemination, we focus on the sharing content in the group and in the respective social connections of its members.

5.8.3 Media Categorization

Our choice of exploiting message categories in the derivation of the traditional collative variables follow prior work [186], including ours [147], which also used preexisting content categories to derive metrics associated with content novelty, uncertainty, conflict and complexity. In our present context, there is no preexisting categorization of content (as exists in platforms like LastFM [147]). Thus, we chose to use a coarse categorization (the only one available in our dataset) – media type. It is a simplification. However, by doing so, we make our metrics independent of any specific method to build such categorization. The multitude of strategies that could be applied to do so – e.g. topic models [59], text embedding strategies [37, 27, 181], exploring a finer grained categorization including memes, emojis, stickers and even more recently deployed WhatsApp features¹⁹ – justifies a study on itself, which should be subject of future work.

Also, as mentioned in Section 5.4.1, a key assumption of ours is that since different media types typically require different levels of engagement of the user to access the content, they should have different effects on curiosity stimulation. For example, textual content is immediately visible, whereas URLs require the user to leave WhatsApp and go to another (Web) application, which is something that the user might feel inclined to defer to another time (or even not do at all). Also, reading a textual message (most often) requires much less effort and much less time than watching a video or listening to an audio. Again, the user may simply choose to defer watching the video or listening the audio to a later time, because she either does not feel like or cannot do it at the time she sees the message. Indeed, this assumption finds some evidence in recent studies of WhatsApp messages which revealed distinct properties in terms of propagation dynamics depending on the type of media [122, 121, 94, 95].

To offer further evidence of the potential distinct effects of different media types in curiosity stimulation, we computed, for each category c ($c \in \{\text{"text"}, \text{"image"}, \text{"audio"}, \text{"video"}, \text{"URL"}\}$), the time interval between the sharing of a message with category c and the next message (regardless of its category) shared in the same group. This time interval is a (rough) estimate of the reaction time, i.e., the time it took for someone in the group to react (by sharing new content) to the first message. Shorter reaction times might suggest a greater impact on user behavior. Obviously, this should be considered only a coarse approximation of reality, as several other (internal and external) factors may have driven the user to share the new content. Yet, we found very different reaction times for the five categories analyzed. Specifically, average reaction times and corresponding 95% confidence intervals, computed for our entire dataset, are as follows: (a) 4.69 ± 0.02

¹⁹<https://faq.whatsapp.com/general/chats/about-forwarding-limits/?lang=en>

minutes for text; (b) 7.73 ± 0.07 minutes for images; (c) 8.10 ± 0.18 minutes for audios; (d) 10.90 ± 0.10 minutes for videos; and (e) 11.61 ± 0.09 minutes for URLs. Clearly, messages with textual content tend to have shorter reaction times, whereas URLs and videos tend to trigger much longer reaction times, which corroborates the assumption discussed in the last paragraph.

5.8.4 Focus on Intra-group Curiosity Stimulation

As a first effort to quantify social curiosity in group communication, we chose to focus on each group separately, under the assumption that *intra-group social influence* is a primary factor driving user curiosity. Indeed, our results suggest that the curiosity of a user may be stimulated differently depending on the group (see, e.g., results in Figures 5.8a and 5.8b as well as Figure 5.9). By focusing on *intra-group* curiosity stimulation we are able to isolate (potentially secondary) *inter-group* effects, offering a foundation upon which follow-up studies can be developed. Moreover, we note that our dataset, as all datasets used in other analyses of WhatsApp [40, 106, 23, 60], are collected from publicly accessible groups. These are essentially different from (private) family and friend groups where members have external (real-world) connections that may influence their curiosity stimulation. In publicly accessible groups, it is not possible to identify such external links among users.

5.8.5 Modeling of Temporal Dynamics

In the derivation of our metrics, we assume that the curiosity driving the sharing of a message has a period of activation δ_T determined by a fixed-duration time window. In other words, we assume that only messages shared within the specific time window stimulate the user’s curiosity. This time window allows us to capture changes in the curiosity stimulation of a user over time. However, it is not possible to determine, from the data, whether indeed the user read all messages shared during a time window. Our assumption is thus a design choice to be able to capture the temporal dynamics of user curiosity.

Such choice, and in particular the window duration of 30 minutes, is based on observations in prior work [22], where authors used the same window duration to study

collective attention in publicly accessible WhatsApp group. The authors experimented with several window durations (from 5 minutes to 2 hours), selecting 30 minutes after manual investigation of the topics discussed in the analyzed groups during the defined windows. They argue that 30-minute windows led to the best tradeoff between capturing small variations of attention over time and still approximating the duration of continuous conversations in WhatsApp. Given the close relationship between attention and curiosity, we chose to follow the same approach, especially because our work focuses on Whatsapp groups similar to those analyzed in [22], that is, publicly accessible political-oriented groups in Brazil.

As an alternative to fixed-duration time windows, one could consider windows with variable durations. For example, one might argue that the window duration should be adjusted based on the current level of activity (message sharing) in the group. However, we found very weak correlations between the number of messages shared during a window and the average user and group-level social curiosity, as estimated by the proposed metrics (Spearman correlation coefficients below 0.14). Nevertheless, investigating alternative approaches to model the temporal dynamics of user curiosity stimulation is an avenue worth pursuing in the future.

5.9 Summary

In this chapter, we have investigated curiosity stimulation as a driving force behind content sharing in group communication, specifically in WhatsApp groups, focusing on a novel component of curiosity stimulation, namely social influence. To that end, we proposed novel metrics to quantify such component, as well as metrics that instantiate other traditionally studied collative variables related to curiosity stimulation (though not to social behavior). We used the proposed metrics to show that social influence is indeed a distinct component of curiosity stimulation, as compared to the other traditional variables, as well as to offer a broad characterization of social curiosity stimulation in a number of publicly accessible political-oriented WhatsApp groups in Brazil.

Recall that, as presented in Section 5.1, we here aimed at addressing three research questions, which we repeat below:

- **RQ2.1:** *How to quantify social influence as a stimulus to one's curiosity driving the information dissemination process?*
- **RQ2.2:** *How does social influence relate to other collative variables priorly associated with curiosity stimulation?*

- **RQ2.3:** *How are users characterized in terms of social stimulation to curiosity?*

Towards answering this RQ2.1, we proposed four new metrics to capture social influence as a component of curiosity stimulation driving individual users to communicate with each other by sharing content in a currently very popular group communication platform – WhatsApp. The derivation of our metrics is founded on the arguments of psychologist Daniel Berlyne on the modeling of human curiosity and its connection with information theory [16].

The challenges we faced while deriving these metrics relate to how to instantiate Berlyne’s general methodology to quantify collative variables associated with curiosity stimulation to the particular context of social influence in group communication. Specifically, we had to make assumptions so as to be able to deal with the lack of explicit signals of user interactions in our dataset. Also, results from prior studies of WhatsApp [22] motivated us to derive metrics capturing not only the traditionally studied direct influence but also indirect influence reflecting the possible presence of some great influencers in the group. Similarly, we argued for the relevance of proposing metrics for both individual users and groups. Notably assessing group-level social curiosity may offer valuable insights into the dynamics of group members provided that groups are reasonably small and discussions are more focused on specific topics, as is the case of WhatsApp groups. This is in contrast to more open spaces of communication (e.g., Twitter, Facebook), where group membership is unlimited and often much larger, and, as such, curiosity stimulation driving the discussions might be more disperse and perhaps more fragmented.

RQ2.2, in turn, relates to how social influence, as captured by the proposed metrics, captures aspects of curiosity stimulation not covered by novelty, complexity, conflict and uncertainty, which are other collative variables associated with curiosity stimulation that have already been studied in several domains (though not in group communication). Towards answering this question, we first presented seven metrics that capture these collative variables for our target environment. The derivation of these metrics is greatly inspired by previous studies of the same variables in other setups, notably our own effort to model user curiosity in LastFM [147], which, in turn, are rooted, once again, on Berlyne’s seminal work.

In particular, as done in prior work, we considered two different elements that may stimulate one’s curiosity, namely: content, captured by the categories associated with the several portions of the message, and the user sharing the content itself. We then proposed metrics that capture the role of both elements on curiosity stimulation with respect to the four aforementioned collative variables.

Having defined these metrics, we answered RQ2.2 by evaluating how correlated the metrics of social influence are to the metrics related to the other variables. This analysis, as all others discussed below, were performed on a large dataset containing over 2 million

messages shared by over 7.5 thousand users in 335 publicly accessible WhatsApp groups in Brazil. The dataset covers a one-year time interval, which includes a period of large use of WhatsApp for the political debate in the country²⁰. As shown in Figure 5.2, our key result is that the four metrics related to social influence are complementary to the other metrics in the vast majority of cases. Thus, social influence, as captured by our proposed metrics, does indeed represent a novel component of curiosity stimulation that is not covered by the other variables. Having addressed RQ2.2, we then proceeded to tackle RQ2.3 by performing a broad characterization of social curiosity stimulation at three levels of aggregation, namely individual messages, individual members of a group (i.e., users) and all members of a group.

We turn to the RQ3 which focuses on the precise modeling of Wundt’s curve, but this time in the context of the production and consumption of information dissemination on the WhatsApp case. To this end, we note that 16% of users are estimated by 2-Beta and 3-Beta mixture models, which is a different percentage than the LastFM case study (Chapter 4). It is clear that these differences relate to user behavior on the different platforms and highlight the existence of the two spectrums of the online information dissemination process, namely information consumption and information production (e.g., content sharing).

In addition, to estimate the curiosity profile (i.e., Wundt’s curve), we use the social curiosity metrics together with the traditional collative variables. We emphasize the contribution of social influence in the composition of the mean stimulus to estimate Wundt’s curve, because by using two sets of variables we reduce the number of invalid users from 33% (when using only social curiosity variables) to 4.8%, a reduction of almost 7-fold, showing the highest accuracy in the estimation of curiosity profiles. This underlines the importance of using these two variables together to estimate the curiosity profile.

The need for the Beta mixture distribution to fit Wundt’s curve arises from the fact that some users switch between different stimulus profiles of curiosity, resulting in multimodal curves of the curiosity profile. If we then consider all Beta distribution models, it appears that the distribution of *Independent* profiles per user is stable regardless of the type of Beta distribution model. In contrast, the distribution of *Dependent* profiles per user shows an increase from a single Beta distribution to a 2-Beta and 3-Beta mixture distribution. Finally, the distribution of *Indirect* profiles per user shows and decrease from 2-Beta to 3-Beta mixture models. The complexity of the users’ curiosity profile is therefore underlined by our results.

If we examine how the curiosity profile evolves over time, this complexity becomes even clearer. To show this, we use the same method as in Chapter 4, i.e. we divide the activity periods of the users into short periods of one week. We thus estimate Wundt’s

²⁰<https://www.niemanlab.org/2018/10/what-to-know-about-whatsapp-in-brazil-ahead-of-sundays-election/>

curve for each week and see that for more than 96% of cases, a 2-Beta or 3-Beta mixture distribution is required to estimate their curiosity profiles. While 76% of cases modeled with the 2-Beta mixture model have at least 2 stimulus profiles and more than half of the users with curiosity profiles modeled with the 3-Beta mixture model compose their stimuli from different stimulus profiles.

As in Chapter 4, we use three different null models in this chapter to test whether our metrics for social curiosity actually capture the social influence of curiosity. Each null model is designed far removed from the user's preferences and therefore closer to a random pattern. So, the differences in the metrics between the real stimulus values and the synthetic stimulus values of the null model are more pronounced depending on the type stimulation profile of social curiosity. For the profiles *Dependent*, for example, the metric *Maximum Direct Influence* dominates over others, then the differences in this metric are greater for this type of stimulus profile. In contrast, for the profile *Indirect* the metric *Maximum Indirect Influence* is dominant and therefore these differences are greater for this metric. Regardless of the null models, the differences for the *Independent* profiles are somewhat smaller, and since this remains the same, they are irrelevant. In this context, we can conclude that these results show that our social curiosity metrics actually capture the stimuli of users' curiosity on the WhatsApp communication platform.

To our knowledge, our findings are novel, with respect to prior efforts to characterize social curiosity as well as prior studies of user behavior in group communication platforms (notably WhatsApp). As such, they offer valuable insights into user behavior (i.e., content sharing) on a platform that has had paramount importance as major source of information (including misinformation) in several countries, with reported impacts on political and social aspects of these societies²¹. Understanding the factors driving users to share content in such platform, social curiosity being one of them, is thus valuable not only from the perspective of human behavior analysis but also because such understanding may motivate future developments to build more socially driven spaces for online communication.

Complementary, our metrics could be applied to analyze social curiosity in other platforms of group communication (e.g., Telegram), or adapted to other types of social media applications (e.g., Twitter, Facebook). Moreover, one could further analyze the relationship between user and group-level social curiosity stimulation. Our proposed social curiosity metrics could be employed, jointly with metrics related to other collative variables, to build more sophisticated curiosity models, which in turn could be explored to improve information dissemination in the target platforms (e.g., as a component of recommendation or search systems).

²¹<https://www.washingtonpost.com/news/theworldpost/wp/2018/11/01/whatsapp-2/>;
<https://www.bbc.com/news/world-asia-india-44856910>;
<https://www.bbc.com/news/education-55332321>

Chapter 6

Conclusions and Future Research Directions

In this chapter we present in Section 6.1 a summary of the results obtained so far and their conclusions, in Section 6.2 we point out future research directions and in Section 6.3 the publications of this dissertation are listed.

6.1 Current Achievements

Recall from Chapter 1 that our hypothesis in this dissertation is that by modeling and analyzing human curiosity we are able to uncover relevant knowledge to better understand information dissemination and to design more effective personalized information services. With that hypothesis as guideline, we have defined four research questions:

- **RQ1:** *Are there distinct user behavior profiles in terms of curiosity stimuli, as captured by multiple collative variables?*
- **RQ2:** *How can we capture social influence as a component of human curiosity stimulation driving online information dissemination?*
- **RQ3:** *To which extent, user curiosity driving online information dissemination can be accurately modeled by a Wundt's curve?*

We refer to the aforementioned questions as *major research questions* (RQs). In contrast, we refer to specific questions raised to investigate each case study as *supplementary research questions*. We illustrate the relationship between the major research questions and supplementary research questions in the diagram presented in Figure 6.1. Major research questions are shown in purple, whereas supplementary questions related to case studies 1 (CS1) and 2 (CS2) are shown in green and salmon, respectively. For




Figure 6.1: Diagram with the major research questions and supplementary research questions associated with each case study.

example, RQ1-CS1 and RQ3-CS1 are supplementary questions for RQ1 and RQ3 respectively, which refer to the first case study of LastFM. RQ2.1-CS2, RQ2.2-CS2, RQ2.3-CS2 and RQ3-CS2 are supplementary questions for RQ2 and RQ3 that relate to our second case study on WhatsApp. More specifically, we recall the supplementary questions for RQ2 of second case study as follows:

- **RQ2.1-CS2:** *How to quantify social influence as a stimulus to one's curiosity driving the information dissemination process?*
- **RQ2.2-CS2:** *How does social influence relate to other collative variables priorly associated with curiosity stimulation?*
- **RQ2.3-CS2:** *How are users characterized in terms of social stimulation to curiosity?*

As shown in the figure (major and supplementary) research questions fall into two complementary layers, to know:

- Curiosity stimulation layer:* encompasses research questions tackling the first step of the general curiosity appraisal process, that is, the derivation of metrics to enable the quantification of different collative variables associated with curiosity stimulation;
- Curiosity score layer:* encompasses research questions related to the second step of the general appraisal process, which aims to relate estimates of stimuli to a curiosity

score driving user behavior. In that particular layer, our focus was on investigating the extent to which a single Wundt's curve (widely used to map curiosity stimuli to curiosity score) is the most adequate model to all users.

Our contributions, summarized in the following subsections, are focused on the two layers, namely curiosity stimulation and curiosity score. Specifically, we have investigated RQ1 and RQ3 as part of case study 1, which is presented in Chapter 4, and also RQ3 as part of case study 2 in Chapter 5. We then turned our attention to RQ2, focused on social influence as a component of curiosity, as we developed case study 2¹, as presented in Chapter 5.

6.1.1 RQ1: Multiple Collative Variables and Discovery of User Curiosity Profiles

The exploration of our first case study (Chapter 4), led to the following main results related to RQ1:

1. We proposed metrics to capture different aspects related to user curiosity stimulation, notably aspects related to the traditionally studied collative variables of uncertainty, novelty, complexity and conflict. We also performed an in depth investigation to identify possible interplay between these variables aiming at identifying complementary as well as redundant metrics. Specifically, we identified that five metrics, namely, novelty (song, artist and musical genre), uncertainty (musical genre) and complexity (musical genre) capture complementary components of curiosity stimulation in most of the cases.
2. Using the five complementary metrics and clustering analysis, we uncovered four different profiles of curiosity stimulation at the access level, namely: (1) *Underground*, characterized by low stimulus levels according to all metrics; (2) *Niche Radical*, characterized by accesses to repeated (low novelty) but complex (multiple musical genres) songs; (3) *Eclectic* refers to an opposite pattern in which curiosity stimulation is driven towards new and different songs; and (4) *Explorer*, which is similar to *Eclectic* but characterized by accesses to more complex songs, i.e., songs with greater diversity of genres.

¹We note that even though we did propose metrics associated with various collative variables in case study 2, the analysis and results are mostly focused on *social influence*.

3. At the user level, by observing the musics a user listens to over time, we were able to uncover dynamic patterns representing how a user’s curiosity stimulation changes across different profiles. In particular, we point out two important patterns. Firstly, there is a clear trend towards repeated behavior, i.e., the user tends to remain on the curiosity stimulation profile over successive listening events. Secondly, though changes in curiosity stimulation profile do occur occasionally, for all profiles, there are larger chances of a user migrating between Eclectic and Explorer profiles;
4. The most of the curiosity stimulation profiles (*Underground*, *Eclectic*, and *Explorer*) show patterns in which users prefer songs with higher stimulus values (in terms of complexity, uncertainty, and novelty), i.e., they prefer new songs by different artists with different musical genres related to those they have recently listened. This suggests that in the process of online information dissemination, users select content item that is different from what they have consumed in the recent past. This is consistent with the way certain online content goes viral for a short period of time (bursts of popularity) and then fades into oblivion (perhaps from boredom due to repeated consumption – familiarity – or loss of attention).

To our knowledge, our attempt to quantify the curiosity stimulation by multiple collative variables as a driver of online music consumption is a novelty that significantly extends the current state of the art. Our observations show that user curiosity stimulation, as captured by the traditionally studied variables, can be quite complex and to some extent highly dynamic. Building on these results, in our next contributions we extend the set of collative variables to include the effects of social influence on curiosity stimulation and explore how curiosity stimuli, as estimated by the proposed metrics, should map to a curiosity score.

6.1.2 RQ2: Social Influence as Component of Human Curiosity Stimulation

Our key take-home messages regarding RQ2, produced as results of our second case study (Chapter 5), can be summarized as:

1. Curiosity stimulation varies greatly not only across individual messages (consistently with observations made in our first case study) but also across different users and different groups (the latter two taken as average stimuli). By employing clustering

analysis, we were able to uncover profiles of social curiosity stimulation at both message and user levels.

2. At the message level, we uncovered three profiles, which indicate that, although social influence does not always play a clear role as a component of curiosity stimulation (profile independent), in almost 30% of the cases, either direct or indirect social influence, as captured by our metrics, are indeed important components driving one's curiosity to share content (profiles dependent and indirect, respectively).
3. At the user level, we showed that the evolution of user curiosity over time with respect to the three message-level stimulation profiles present great heterogeneity, revealing five different user profiles. While some users do exhibit a rather stable curiosity stimulation process, remaining at the same message-level profile over time, for many other users, the role of (direct/indirect) social influence as a component of curiosity stimulation changes greatly over time, probably in response to the dynamics of interactions within the group.
4. We also showed evidence that curiosity stimulation may change significantly depending on the group the user participates in. This observation validates one of our key assumptions (Chapter 3), offering also insights into the role that such (often small) groups have as drivers to user behavior (notably content sharing). In some sense, these results mimic what we see in real life when the same person may behave quite differently (more or less participative/chatty) depending on the group of people she is interacting with.
5. At the group level, we noted that the role of social curiosity is more clearly observed when group members are more engaged in the ongoing discussions, more often sharing content (thus interacting with others), as these actions suggest a greater susceptibility to social curiosity.

6.1.3 RQ3: Investigating Wundt's Curve as Model of Human Curiosity

In a second step developed in our two case studies, we explored the proposed metrics of curiosity stimulation to investigate the suitability of Wundt's curve as a model for curiosity elicitation. Our main findings revealed that:

1. The traditionally used simple bell-shaped Wundt curve does not adequately model the process of arousing the curiosity of a large portion (almost half) of the users

in the first case study, namely the LastFM platform, and a representative portion (showing the presence of behavior that deviates from the majority) of the users in the second case study, i.e. the communication platform WhatsApp. For these users, we found instead that the curiosity curve has a multimodal shape that may reflect the quite complex and dynamic patterns we uncovered in our analyzes of the curiosity stimulation metrics.

2. For such users, we suggest using a combination of two or three Beta distributions, as these can capture the different distribution modes and thus reduce errors in identifying the areas of greater curiosity stimulation.

Furthermore, these observations can of course be extended to domain other than music and communication platforms (e.g., WhatsApp and Telegram), such as content production and sharing services (e.g., Twitter, Instagram and TikTok). Finally, these findings form the basis for designing more effective personalized services.

6.2 Future Research Directions

In this section we discuss future research directions that we intend to pursue. These research directions are thus divided into two possibilities of research as a form of extension of our main findings.

6.2.1 Application of Curiosity Models in the Information Dissemination

In examining the hypothesis of this dissertation, our results have shown that an important contribution is the fact that social influence is a component to stimulate curiosity in the process of online information dissemination, especially in social media applications where users are influenced by their social connections. However, the role of social curiosity in the process of information dissemination can be better explored when the representation of content is available in the dataset. In the course of our methodology for developing curiosity metrics in the two case studies, we have shown that the categorization of content items (i.e., media) forms the basis for capturing the stimuli for human curiosity. Thus,

by accessing the content of users' messages, it is possible to investigate how content influences social curiosity to drive the process of information dissemination. So, in this sense, it is possible to create a fine-grained media categorization based on content properties of audio, video, text, URLs, and images.

In the case of text messages, for example, advanced language models [35] can be used to topic modeling for text categorization to enable a more comprehensive analysis of social curiosity stimulus. Besides, the use of modern machine learning techniques to captioning and labeling media (i.e., image [41], video [182] and audio [180]) can help to find specific categorizations for these types of media and thus identify their impact on social curiosity. In addition, it is important to point out that we are aware that there are other aspects that lead users to share information in the process of information dissemination, such as polarization, opinion formation and belief formation. With access to media (video, image and audio) and text messages, as well as advanced machine learning techniques and algorithms, these aspects, which are not addressed in this dissertation, can be explored in future work.

As mentioned in Chapter 2 on the social influence background, there are two important types of social influence in information dissemination process: (i) *quantification of social influence* and (ii) *information diffusion models*. Our contribution focused on the former type, but we recognise the importance of the latter. Among these information diffusion models can be cited epidemic models, cascade models, linear threshold models, and complex contagion models [103]. To our knowledge, Vega-Oliveros *et al.* were the first authors to include social curiosity in an epidemic model [161]. Beyond this, we are not aware of any other work dealing with social curiosity and information diffusion models, so this is a research gap that needs to be filled and it as a possible future research direction. Finally, complex contagion models also differ from the other models mentioned here, as the success of the information dissemination process depends on the user being exposed to multiple social peers before adopting a new behavior [50]. In this sense, it is important to define as future work a research that investigates in depth how social curiosity influences complex contagion models in different types of online social platforms.

6.2.2 Application of Curiosity Models in Recommendation Systems

For the second future research direction, we propose a hybrid framework of recommendations that should include the following:

- (a) *General appraisal process*: this consists of a complete set of collative variables (including social curiosity) and our Beta-mixture to model Wundt’s curve for heterogeneous and dynamic behavior of online users. Our intuition is that by introducing the general appraisal process of curiosity into the recommendation machinery, we will be able to create recommendations that better match the idiosyncrasies of individual users. To achieve this, we intend to include metrics related to all the collative variables considered in the two case studies, i.e., the traditional collative variables and the social variables. To this end, we plan to take the following steps: (1) adapt the social curiosity metric to an information consumption scenario that corresponds to the recommender systems domain, to do this, we can explore the user-rating matrix and the explicit social connections that may exist, as was done by [174, 170]; (2) examine the set of candidate metrics that capture complementary aspects of curiosity stimulation by applying the methodology to identify a reduced set of selected metrics defined in this dissertation; and (3) incorporate heterogeneous and multi-modal models that map curiosity stimuli to curiosity values by using a Beta-mixture distribution as a model of Wundt’s curve.
- (b) *A machine learning-based recommendation model*: which takes into account a range of content and user attributes, including elements of curiosity. This machine learning-based recommendation model should be based on *the artificial neural networks*, as these are generally able to uncover complex and non-linear relationships between users and items [3]. Moreover, online social network data exhibit non-linear relationships that influence information dissemination [97, 34]. Nevertheless, recent studies have already used *neural networks* to solve sparsity and cold-start problems in recommendation algorithms [53, 67]. All these characteristics making *neural networks* an excellent option for building hybrid recommendation systems.

Finally, we intend to evaluate the performance of our proposed framework against the main state-of-the-art baselines selected from the literature, which are summarized in Table 2.1 (Section 2.2.2). The next step is to evaluate to what extent (and possibly in which scenarios) the addition of the curiosity models improves the quality of the recommendations. To this end, as in the work of Belém *et al.*, it is important to find a sweet spot between the *classical trade-off between accuracy and diversity* [14], where an improvement in diversity leads to a decrease in accuracy and vice versa [11].

6.3 Publications

The results we obtained so far, covering each case study, are summarized in the following publications:

- Sousa, A. M.; Almeida, J. M.; Figueiredo, F.. Analyzing and Modeling User Curiosity in Online Content Consumption: A LastFM Case Study. IEEE/ACM The International Conference Series on Advances in Social Network Analysis and Mining (ASONAM), 2019.
- Sousa, A. M.; Almeida, J. M.; Figueiredo, F.. Metrics of Social Curiosity: The WhatsApp Case. Online Social Networks and Media (OSNEM), 2022.

Bibliography

- [1] Fakhri Abbas and Xi Niu. One size does not fit all: Modeling users' personal curiosity in recommender systems. *arXiv preprint*, arXiv:1907.00119, 2019.
- [2] Lada A Adamic and Eytan Adar. Friends and neighbors on the web. *Social Networks*, 25(3), 2003.
- [3] M. Mehdi Afsar, Trafford Crump, and Behrouz Far. Reinforcement learning based recommender systems: A survey. *ACM Comput. Surv.*, jun 2022. Just Accepted.
- [4] Mehran Ahmadlou, Janou H. W. Houba, Jacqueline F. M. van Vierbergen, Maria Giannouli, Geoffrey-Alexander Gimenez, Christiaan van Weeghel, Maryam Darbanfouladi, Maryam Yasamin Shirazi, Julia Dziubek, Mejdy Kacem, Fred de Winter, and J. Alexander Heimel. A cell type-specific cortico-subcortical brain circuit for investigatory and novelty-seeking behavior. *Science*, 372(6543), 2021.
- [5] I. F. Akyildiz and et al. 5g roadmap: 10 key enabling technologies. *Computer Networks*, 106:17–48, 2016.
- [6] Ahmad Al-Doulat. Surprise and curiosity in a recommender system. In *2018 IEEE/ACS 15th International Conference on Computer Systems and Applications (AICCSA)*, pages 1–2. IEEE, 2018.
- [7] Helena Alicart, David Cucurell, and Josep Marco-Pallarés. Gossip information increases reward-related oscillatory activity. *NeuroImage*, 210:116520, 2020.
- [8] Renato Cordeiro Amorim and Christian Hennig. Recovering the number of clusters in data sets with noise features using feature rescaling factors. *Information Sciences*, 324, 2015.
- [9] M. Arif, C. N. Faisal, H. Ahmad, M. A. Habib, M. Ahmad, and N. Mehmood. The moderating role of curiosity between interactivity and cognitive motives. In *2020 IEEE 23rd International Multitopic Conference (INMIC)*, pages 1–6. IEEE, 2020.
- [10] M. Arnone and R. V. Small. Arousing and sustaining curiosity: Lessons of arcs model. In *Proceedings of the Annual National Conference of the Association of Educational Communications and Technology (AECT)*, pages 1–17, Anaheim, CA, USA, 1995. Association of Educational Communications and Technology (AECT).

- [11] Iman Avazpour, Teerat Pitakrat, Lars Grunske, and John Grundy. Dimensions and metrics for evaluating recommendation systems. In Martin P. Robillard, Walid Maalej, Robert J. Walker, and Thomas Zimmermann, editors, *Recommendation Systems in Software Engineering*, pages 245–273. Springer, 2014.
- [12] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web, WWW '12*, page 519–528, New York, NY, USA, 2012. Association for Computing Machinery (ACM).
- [13] Sarah Bechtle, Yixin Lin, Akshara Rai, Ludovic Righetti, and Franziska Meier. Curious ilqr: Resolving uncertainty in model-based rl. In *Proceedings of the 3rd Conference on Robot Learning*, volume 100, pages 162–171, Osaka, Japan, 2019. PMLR.
- [14] Fabiano M. Belém, Carolina S. Batista, Rodrygo L. T. Santos, Jussara M. Almeida, and Marcos A. Gonçalves. Beyond relevance: Explicitly promoting novelty and diversity in tag recommendation. *ACM Transactions on Intelligent Systems and Technology*, 7(3), 2016.
- [15] Austin R. Benson, Ravi Kumar, and Andrew Tomkins. Modeling user consumption sequences. In *Proceedings of the 25th international conference on World Wide Web, WWW'16*, pages 519–529, Republic and Canton of Geneva, CHE, 2016. International World Wide Web Conferences Steering Committee.
- [16] D.E. Berlyne. *Conflict, Arousal and Curiosity*. McGraw-Hill series in psychology. McGraw-Hill, NY, USA, 1960.
- [17] Sheng Bin, Chi-Cheng Chen, and Gengxin Sun. Maximizing social influence in nearly optimal time: Sris model. In *2020 IEEE 2nd Eurasia Conference on Biomedical Engineering, Healthcare and Sustainability (ECBIOS)*, pages 201–203. IEEE, 2020.
- [18] Francesco Bonchi, Francesco Gullo, Bud Mishra, and Daniele Ramazzotti. Probabilistic causal analysis of social influence. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, pages 1003–1012, New York, NY, USA, 2018. Association for Computing Machinery (ACM).
- [19] Terry Bossomaier, Lionel Barnett, Michael Harré, and Joseph T. Lizier. *Information Theory*. Springer International Publishing, Cham, 2016.
- [20] Gerlof Bouma. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL*, 30:31–40, 2009.

- [21] Gregory J. Boyle. Breadth-depth or state-trait curiosity? a factor analysis of state-trait curiosity and state anxiety scales. *Personality and Individual Differences*, 10(2):175–183, 1989.
- [22] J. A. Caetano, J. M. Almeida, M. A. Gonçalves, Wagner Meira Jr., H. T. Marques-Neto, and V. A. F. Almeida. Analyzing topic attention in online small groups. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '21*, pages 1–5. Association for Computing Machinery (ACM), 2021.
- [23] J. A. Caetano, G. Magno, M. A. Gonçalves, J. M. Almeida, H. T. Marques-Neto, and V. A. F. Almeida. Characterizing attention cascades in whatsapp groups. In *Proceedings of the 10th ACM Conference on Web Science, WebSci '19*, page 27–36. Association for Computing Machinery (ACM), 2019.
- [24] Jessy Ceha, Nalin Chhibber, Joslin Goh, Corina McDonald, Pierre-Yves Oudeyer, Dana Kulić, and Edith Law. Expression of curiosity in social robots: Design, perception, and effects on behaviour. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, pages 1–12, New York, NY, USA, 2019. ACM.
- [25] Damon Centola. Influential networks. *Nature Human Behaviour*, 3(1):664–665, 2019.
- [26] Damon Centola and Michael Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–734, 2007.
- [27] Wei-Cheng Chang, Hsiang-Fu Yu, Kai Zhong, Yiming Yang, and Inderjit S. Dhillon. Taming pretrained transformers for extreme multi-label text classification. In *Proc. of the 26th ACM SIGKDD, KDD '20*, pages 3163–3171, New York, NY, USA, 2020. Association for Computing Machinery.
- [28] Bo-Lun Chen, An Zeng, and Ling Chen. The effect of heterogeneous dynamics of online users on information filtering. *Physics Letters A*, 379(43):2839–2844, 2015.
- [29] Li Chen, Yonghua Yang, Ningxia Wang, Keping Yang, and Quan Yuan. How serendipity improves user satisfaction with recommendations? a large-scale user evaluation. In *The World Wide Web Conference, WWW' 19*, page 240–250, New York, NY, USA, 2019. Association for Computing Machinery (ACM).
- [30] Robert P. Collins, Jordan A. Litman, and Charles D. Spielberger. The measurement of perceptual curiosity. *Personality and Individual Differences*, 36(5):1127–1141, 2004.

- [31] Federico Coró, Gianlorenzo D’angelo, and Yllka Velaj. Link recommendation for social influence maximization. *ACM Transactions on Knowledge Discovery from Data*, 15(6):94:1–94:23, May 2021.
- [32] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, NY, USA, 2006.
- [33] Easley David and Kleinberg Jon. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, USA, 2010.
- [34] Easley David and Kleinberg Jon. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, USA, 2010.
- [35] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, Minneapolis, Minnesota, jun 2019. Association for Computational Linguistics.
- [36] Malcolm Doering, Phoebe Liu, Dylan F. Glas, Takayuki Kanda, Dana Kulić, and Hiroshi Ishiguro. Curiosity did not kill the robot: A curiosity-based learning system for a shopkeeper robot. *ACM Transactions Human-Robot Interaction*, 8(3):1–24, jul 2019.
- [37] Ramazan Esmeli, Mohamed Bader-El-Den, and Hassana Abdullahi. Using word2vec recommendation for improved purchase prediction. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2020.
- [38] Carlos H. G. Ferreira, Fabricio Murai, Breno de Souza Matos, and Jussara M. Almeida. Modeling dynamic ideological behavior in political networks. *The Journal of Web Science*, 7:1–14, 2019.
- [39] Zhe Fu and Xi Niu. Modeling users’ curiosity in recommender systems. *ACM Transactions Knowledge Discovery and Data Mining*, 18(1), 2023.
- [40] Kiran Garimella and Gareth Tyson. Whatsapp doc?: A first look at whatsapp public group data. *ICWSM*, 2018.
- [41] Taraneh Ghandi, Hamidreza Pourreza, and Hamidreza Mahyar. Deep learning approaches on image captioning: A review. *ACM Computing. Surveys*, 56(3), oct 2023.
- [42] R. Golman and G. Loewenstein. Information gaps: A theory of preferences regarding the presence and absence of information. *Decision - American Psychological Association*, 5(3):143–164, 2016.

- [43] R. Golman and G Loewenstein. The desire for knowledge and wisdom. In Gore Gordon, editor, *The New Science of Curiosity*, chapter 1, pages 1–5. Nova Science Publishers, 2018.
- [44] G. Gordon, C. Breazeal, and S. Engel. Can children catch curiosity from a social robot? In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '15*, pages 91–98, New York, NY, USA, 2015. ACM/IEEE.
- [45] J. Gottlieb and et al. Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11):585–593, 2013.
- [46] Amit Goyal, Francesco Bonchi, and Laks V.S. Lakshmanan. Learning influence probabilities in social networks. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM '10*, pages 241–250, New York, NY, USA, 2010. Association for Computing Machinery (ACM).
- [47] Kazjon Grace and Mary Lou Maher. Specific curiosity as a cause and consequence of transformational creativity. In Hannu Toivonen, Simon Colton, Michael Cook, and Dan Ventura, editors, *Proceedings of the Sixth International Conference on Computational Creativity, ICCCI 2015*, pages 260–267, Park City, Utah, USA, 2015. ICCCI.
- [48] A. Greasley, A. Lamont, and J. Sloboda. Exploring musical preferences: An in-depth qualitative study of adults’ liking for music in their personal collections. *Qualitative Research in Psychology*, 10(4), 2013.
- [49] Douglas Guilbeault, Joshua Becker, and Damon Centola. *Complex Contagions: A Decade in Review*, chapter Part I: Introduction to Spreading in Social Systems, pages 3–25. Springer International Publishing, Cham, 2018.
- [50] Douglas Guilbeault and Damon Centola. Topological measures for identifying and predicting the spread of complex contagions. *Nature Communications*, 12(1):4430, 2021.
- [51] Yuning Guo, Jianxiang Cao, and Weiguo Lin. Social network influence analysis. In *2019 6th International Conference on Dependable Systems and Their Applications (DSA)*, pages 517–518, Harbin, China, 2019. IEEE.
- [52] F. M. Hartung and B. Renner. Social curiosity and gossip: Related but different drives of social functioning. *PLoS One*, 8(7):e69996, 2013.
- [53] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceed. of the 26th WWW, WWW' 17*, pages

- 173–182, Republic and Canton of Geneva, CHE, 2017. International World Wide Web Conferences Steering Committee.
- [54] Christopher K Hsee and Bowen Ruan. The pandora effect: The power and peril of curiosity. *Psychological science*, 27(5):659–666, May 2016.
- [55] Hui-Ju Hung, De-Nian Yang, and Wang-Chien Lee. Social influence-aware reverse nearest neighbor search. *ACM Transactions on Spatial Algorithms and Systems*, 2(3), oct 2016.
- [56] Ruben Interian, Rusl  n G. Marzo, Isela Mendoza, and Celso C. Ribeiro. Network polarization, filter bubbles, and echo chambers: an annotated review of measures and reduction methods. *International Transactions in Operational Research*, 30(6):3122–3158, 2023.
- [57] Sergei Ivanov, Konstantinos Theocharidis, Manolis Terrovitis, and Panagiotis Karas. Content recommendation for viral social influence. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’17, pages 565–574, New York, NY, USA, 2017. Association for Computing Machinery (ACM).
- [58] Oliver P. John and Sanjay Srivastava. The big five trait taxonomy: History, measurement, and theoretical perspectives. In *Handbook of personality: Theory and research*, volume 2, pages 102–138, New York, USA, 1999. Guilford Press.
- [59] Dakshi Tharanga Kapugama Geeganage. Concept embedded topic modeling technique. In *Companion Proceedings of the The Web Conference 2018*, page 831–835, 2018.
- [60] P. Kariuki and L. O. Ofusori. Whatsapp-operated stokvels promoting youth entrepreneurship in durban, south africa: Experiences of young entrepreneurs. In *10th ICEGOV*, 2017.
- [61] Todd B. Kashdan, David J. Disabato, Fallon R. Goodman, and Patrick E. McKnight. The five-dimensional curiosity scale revised (5dcr): Briefer subscales while separating overt and covert social curiosity. *Personality and Individual Differences*, 157:109836, 2020.
- [62] Todd B. Kashdan, Fallon R. Goodman, David J. Disabato, Patrick E. McKnight, Kerry Kelso, and Carl Naughton. Curiosity has comprehensive benefits in the workplace: Developing and validating a multidimensional workplace curiosity scale in united states and german employees. *Personality and Individual Differences*, 155:109717, 2020.

- [63] Todd B. Kashdan, Melissa C. Stikma, David J. Disabato, Patrick E. McKnight, John Bekier c, Joel Kaji, and Rachel Lazarus. The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality*, 73:130–149, 2018.
- [64] Leonard Kaufman and Peter J. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley, USA, 2005.
- [65] C. Kidd and B. Y. Hayden. The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460, November 2015.
- [66] Isabel Kloumann, Lada Adamic, Jon Kleinberg, and Shaomei Wu. The lifecycles of apps in a social ecosystem. In *Proc. WWW*, 2015.
- [67] Hyeyoung Ko, Suyeon Lee, Yoonseo Park, and Anna Choi. A survey of recommendation systems: Recommendation models, techniques, and application fields. *Electronics*, 11(1), 2022.
- [68] M. S. Kosinski. *Measurement and Prediction of Individual and Group Differences in The Digital Environment*. PhD thesis, U. Cambridge, UK, 2014.
- [69] Arun Kumar and Paul Schrater. Novelty learning via collaborative proximity filtering. In *Proceedings of the 22nd IUI*, 2017.
- [70] Nikhil Kumar, Ruocheng Guo, Ashkan Aleali, and Paulo Shakarian. An empirical evaluation of social influence metrics. In *2016 IEEE/ACM ASONAM*, ASONAM’16, pages 1329–1336, Davis, California, USA, 2016. IEEE Press - ASONAM.
- [71] Renaud Lambiotte and Michal Kosinski. Tracking the digital footprints of personality. *Proceedings of the IEEE*, 102(12):1934–1939, 2014.
- [72] Johnny King L. Lau, Hiroki Ozono, Kei Kuratomi, Asuka Komiya, and Kou Murayama. Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks. *Nature Human Behaviour*, 4:531–543, 2020.
- [73] E. Law and et al. Curiosity killed the cat, but makes crowdwork better. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI’16, page 4098–4110, NY, USA, 2016. Association for Computing Machinery (ACM).
- [74] Edith Law, Parastoo Baghaei Ravari, Nalin Chhibber, Dana Kulic, Stephanie Lin, Kevin D. Pantasdo, Jessy Ceha, Sangho Suh, and Nicole Dillen. Curiosity notebook: A platform for learning by teaching conversational agents. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI EA’20, pages 1–9, New York, NY, USA, 2020. Association for Computing Machinery.

- [75] Boyao Li, Tao Lu, Jiayi Li, Ning Lu, Yinghao Cai, and Shuo Wang. Acder: Augmented curiosity-driven experience replay. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4218–4224. IEEE, 2020.
- [76] Chunfang Li and Fei Xiong. Social recommendation with multiple influence from direct user interactions. *IEEE Access*, 5:16288–16296, 2017.
- [77] Kan Li, Lin Zhang, and Heyan Huang. Social influence analysis: Models, methods, and evaluation. *Engineering*, 4(1):40–46, 2018. Cybersecurity.
- [78] Lin Li, Ang Li, Bibo Hao, Zengda Guan, and Tingshao Zhu. Predicting active users’ personality based on micro-blogging behaviors. *PLOS ONE*, 9(1):1–11, 01 2014.
- [79] Ye Li, Hong Xie, Yishi Lin, and John C.S. Lui. To be or not to be: Analyzing amp; modeling social recommendation in online social networks. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 1180–1185. IEEE, 2019.
- [80] Jordan A. Litman and Mark V. Pezzo. Dimensionality of interpersonal curiosity. *Personality and Individual Differences*, 43(6):1448–1459, 2007.
- [81] Siyuan Liu, Shuhui Wang, and Feida Zhu. Structured learning from heterogeneous behavior for social identity linkage. *IEEE Transactions on Knowledge and Data Engineering*, 27(7):2005–2019, 2015.
- [82] M. Livio. *Why?: What Makes Us Curious*. Symon & Schuster, NY, USA, 2017.
- [83] G. Loewenstein. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1):75–98, 1994.
- [84] G. Loewenstein. Recommender systems and the new new economics of information. In *Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys ’17*, page 1, New York, NY, USA, 2017. Association for Computing Machinery (ACM).
- [85] Alvis Logins and Panagiotis Karras. Content-based network influence probabilities: Extraction and application. In *2019 International Conference on Data Mining Workshops (ICDMW)*, pages 69–72. IEEE, 2019.
- [86] Luca Luceri, Torsten Braun, and Silvia Giordano. Analyzing and inferring human real-life behavior through online social networks with social influence deep learning. *Applied Network Science*, pages 1–25, 2019.
- [87] Yadan Luo, Zi Huang, Zheng Zhang, Ziwei Wang, Jingjing Li, and Yang Yang. Curiosity-driven reinforcement learning for diverse visual paragraph generation. In *Proceedings of the 27th ACM International Conference on Multimedia, MM ’19*,

- pages 2341–2350, New York, NY, USA, 2019. Association for Computing Machinery (ACM).
- [88] V. Maccatrozzo and et al. Everybody, more or less, likes serendipity. In *Proceedings 25th ACM UMAP*, UMAP '17, pages 29–34, New York, NY, USA, 2017. Association for Computing Machinery (ACM).
 - [89] V. Maccatrozzo and et al. Sirup: Serendipity in recommendations via user perceptions. In *Proceedings of the 22nd ACM IUI*, 2017.
 - [90] L. Macedo and A. Cardoso. The role of surprise, curiosity and hunger on exploration of unknown environments populated with entities. In *2005 Portuguese Conference on Artificial Intelligence (EPIA)*, pages 47–53, Covilha, Portugal, 2005. IEEE.
 - [91] Charles Mackay. *Extraordinary popular delusions and the madness of crowds*. Simon and Schuster, 2012.
 - [92] D. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, Cambridge, UK, 2005.
 - [93] Mary Lou Maher, Kathryn E. Merrick, and R. Saunders. Achieving creative behaviour using curious learning agents. In *Proceedings of AAAI Spring Symposium on Creative Intelligent Systems*, pages 40–46, USA, 01 2008. AAAI.
 - [94] Alexandre Maros, Jussara Almeida, Fabrício Benevenuto, and Marisa Vasconcelos. Analyzing the use of audio messages in whatsapp groups. In *Proceedings of the ACM WWW*, 2020.
 - [95] Alexandre Maros, Jussara M. Almeida, and Marisa Vasconcelos. A study of misinformation in audio messages shared in whatsapp groups. In Jonathan Bright, Anastasia Giachanou, Viktoria Spaiser, Francesca Spezzano, Anna George, and Alexandra Pavliuc, editors, *Disinformation in Open Online Media*, pages 85–100, Cham, 2021. Springer International Publishing.
 - [96] Caroline B Marvin and Daphna Shohamy. Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology: General*, 145(3):266–272, March 2016.
 - [97] Yasuko Matsubara, Yasushi Sakurai, B. Aditya Prakash, Lei Li, and Christos Faloutsos. Nonlinear dynamics of information diffusion in social networks. *ACM Transactions on the Web*, 11(2):1–40, April 2017.
 - [98] Philipe F. Melo, Johnnatan Messias, Gustavo Resende, Kiran Garimella, Jussara M. Almeida, and Fabrício Benevenuto. Whatsapp monitor: A fact-checking system for whatsapp. In *Proc. ICWSM*, 2019.

- [99] Daniel A. Menasce and Virgilio Almeida. *Scaling for E Business: Technologies, Models, Performance, and Capacity Planning*. Pren. Hall, USA, 2000.
- [100] Satya Menon and Dilip Soman. Managing the power of curiosity for effective web advertising strategies. *Journal of Advertising*, 31(3):1–14, 2002.
- [101] Kathryn E. Merrick and Mary Lou Maher. *Motivated Reinforcement Learning: Curious Characteres for Multiuser Gamers*. Springer Berlin, Heidelberg, Germany, 1 edition, 2009.
- [102] J. D. Millan-Cifuentes and et al. Curiosity driven search: When is relevance irrelevant? In *Proceedings of the 5th Information Interaction in Context Symposium, IliX 2014, IliX '14*, pages 279–282, NY, USA, 2014. Association for Computing Machinery (ACM).
- [103] Byungjoon Min and Maxi San Miguel. Competing contagion processes: Complex contagion triggered by simple contagion. *Scientific Reports*, 8(1):10422, 2018.
- [104] Maryam Mohseni, Mary Lou Maher, Kazjon Grace, Nadia Najjar, Fakhri Abbas, and Omar Eltayeb. Pique: Recommending a personalized sequence of research papers to engage student curiosity. In *Artificial Intelligence in Education*, volume 11626, pages 201–205, Cham, 2019. Springer International Publishing.
- [105] A. C. B. L. Monção, S. L. Correa, A. C. Viana, and V. K. Cardoso. Optimizing Content Caching and Recommendations with Context Information in Multi-Access Edge Computing. Research report, INRIA Saclay - Ile de France (INRIA), Aug 2021.
- [106] Andrés Moreno, Philip Garrison, and Karthik Bhat. Whatsapp for monitoring and response during critical events: Aggie in the ghana 2016 election. In *14th Information Systems for Crisis Response And Management*, 2017.
- [107] Mehdi Moussaïd, Juliane E. Kämmer, Pantelis P. Analytis, and Hansjörg Neth. Social influence and the collective dynamics of opinion formation. *PLOS ONE*, 8(11):1–8, 11 2013.
- [108] Huda Nassar, Austin R. Benson, and David F. Gleich. Pairwise link prediction. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '19*, pages 386–393, New York, NY, USA, 2019. Association for Computing Machinery (ACM).
- [109] Sumit Negi and Santanu Chaudhury. Link prediction in heterogeneous social networks. In *Proceedings of the 25th ACM International on Conference on Information*

- and Knowledge Management*, CIKM '16, pages 609–617, New York, NY, USA, 2016. Association for Computing Machinery.
- [110] Esther Niehoff and Suzanne Oosterwijk. To know, to feel, to share? exploring the motives that drive curiosity for negative content. *Current Opinion in Behavioral Sciences*, 35, 2020.
- [111] Xi Niu and Ahmad Al-Doulat. Luckyfind: Leveraging surprise to improve user satisfaction and inspire curiosity in a recommender system. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*, CHIIR '21, pages 163–172, New York, NY, USA, 2021. Association for Computing Machinery.
- [112] Gabriel Peres Nobre, Carlos Henrique Gomes Ferreira, and Jussara Marques de Almeida. Beyond groups: Uncovering dynamic communities on the whatsapp network of information dissemination. In *12th Social Informatics*, 2020.
- [113] Gabriel Peres Nobre, Carlos H.G. Ferreira, and Jussara M. Almeida. A hierarchical network-oriented analysis of user participation in misinformation spread on whatsapp. *Information Processing & Management*, 59(1):102757, 2022.
- [114] William Odom, Ron Wakkary, Jeroen Hol, Bram Naus, Pepijn Verburg, Tal Amram, and Amy Yo Sue Chen. Investigating slowness as a frame to design longer-term experiences with personal data: A field study of olly. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI'19, pages 1–16, New York, NY, USA, 2019. ACM.
- [115] R. L. Ott and M. T. Longnecker. *An Introduction to Statistical Methods and Data Analysis*. Cengage Learning, USA, 7th edition, 2015.
- [116] P.-Y. Oudeyer, J. Gottlieb, and M. Lopes. *Intrinsic Motivation, Curiosity, and Learning: Theory and Applications in Educational Technologies*, volume 229 of *Progress in Brain Research*, chapter 11, pages 257–284. Elsevier, 2016.
- [117] Beatrice Perez, Mirco Musolesi, and Gianluca Stringhini. You are your metadata: Identification and obfuscation of social media users using metadata information. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 1–10, New Orleans, Louisiana, USA, 03 2018. Association for the Advancement of Artificial Intelligence.
- [118] A. M. Ramos, N. Andrade, and L. B. Marinho. Exploring the relation between novelty aspects and preferences in music listening. In *Proceedings of the 14th ISMIR*, 2013.

- [119] Neha Rani, Sharon Lynn Chu, Yvette G. Williamson, and Sindy Wu. *Curiosity-Inspired Learning: Insitu versus Post-Event Approaches to Recall and Reflection*, pages 1–6. CHI EA '21. Association for Computing Machinery (ACM), New York, NY, USA, 2021.
- [120] Britta Renner. Curiosity about people: The development of a social curiosity measure in adults. *Journal of Personality Assessment*, 87(3):305–316, 2006.
- [121] Gustavo Resende, Philipe Melo, Julio C. S. Reis, Marisa Vasconcelos, Jussara M. Almeida, and Fabrício Benevenuto. Analyzing textual (mis)information shared in whatsapp groups. In *10th ACM Conference on Web Science*, 2019.
- [122] Gustavo Resende, Philipe Melo, Hugo Sousa, Johnnatan Messias, Marisa Vasconcelos, Jussara Almeida, and Fabrício Benevenuto. (mis)information dissemination in whatsapp: Gathering, analyzing and countermeasures. In *Proc. WWW, WWW '19*, pages 818–828, New York, NY, USA, 2019. Association for Computing Machinery.
- [123] Ana Carolina Ribeiro, Bruno Azevedo, Jorge Oliveira e Sá, and Ana Alice Baptista. How to measure influence in social networks? In Fabiano Dalpiaz, Jelena Zdravkovic, and Pericles Loucopoulos, editors, *Research Challenges in Information Science*, pages 38–57, Cham, 2020. Springer International Publishing.
- [124] Sovan Samanta, Vivek Kumar Dubey, and Biswajit Sarkar. Measure of influences in social networks. *Applied Soft Computing*, 99:106858, 2021.
- [125] Alan Menk dos Santos. A hybrid recommendation system based on human curiosity. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 367–370, New York, NY, USA, 2015. Association for Computing Machinery (ACM).
- [126] Alan Menk dos Santos and L. Sebastiá. Predicting the human curiosity from users' profiles on facebook. In *Proceedings of the 4th Spanish Conference on Information Retrieval, CERI '16*, pages 1–8, NY, USA, 2016. Association for Computing Machinery (ACM).
- [127] Alan Menk dos Santos, Laura Sebastia, and Rebeca Ferreira. Curumim: A serendipitous recommender system based on human curiosity. *Procedia Computer Science*, 112:484–493, 2017.
- [128] R. J. D. F. Santos. Inferring the curiosity by using facebook profile data. Master's thesis, Department of Information Systems and Computation, Univ. Polit. València, Spain, September 2017.
- [129] R. Saunders and J. S. Gero. How to study artificial creativity. In *Proceedings of the 4th ACM Creativity & Cognition, C&C '02*, pages 80–87, New York, NY, USA, 2002. Association for Computing Machinery (ACM).

- [130] M. Schaekermann and et al. Curiously motivated: Profiling curiosity with self-reports and behaviour metrics in the game “destiny”. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, CHI PLAY ’17, page 143–156, New York, NY, USA, 2017. Association for Computing Machinery (ACM).
- [131] M. Schedl. The lfm-1b dataset for music retrieval and recommendation. In *Proc. ACM ICMR*, 2016.
- [132] M. Schedl and B. Ferwerda. Large-scale analysis of group-specific music genre taste from collaborative tags. In *Proc. IEEE ISM*, 2017.
- [133] Markus Schedl. Genre differences of song lyrics and artist wikis: An analysis of popularity, length, repetitiveness, and readability. In *The World Wide Web Conference*, WWW’19, pages 3201–3207, New York, NY, USA, 2019. ACM.
- [134] Andreas Schneider, Georg von Krogh, and Peter Jäger. “what’s coming next?” epistemic curiosity and lurking behavior in online communities. *Computers in Human Behavior*, 29(1):293–303, 2013.
- [135] Grant Schoenebeck and Biaoshuai Tao. Influence maximization on undirected graphs: Toward closing the $(1-1/e)$ gap. *ACM Transactions on Economics and Computation*, 8(4), nov 2020.
- [136] T. Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000.
- [137] C. Schröder and S. Rahmann. A hybrid parameter estimation algorithm for beta mixtures and applications to methylation state classification. *Algorithms for Molecular Biology*, 12(1), 2017.
- [138] D. Sculley. Web-scale k-means clustering. In *Proceedings of the ACM WWW*, 2010.
- [139] Tushar Shandhilya and Nisheeth Srivastava. Using conceptual incongruity as a basis for making recommendations. In *Fourteenth ACM Conference on Recommender Systems*, 2020.
- [140] Arjun Shankar. Capitalism, curiosity, and specialization. *The Newsletter of the Technical Committee on Cognitive and Development Systems*, 15(1):12, 2018.
- [141] Jyoti Shokeen and Chhavi Rana. A study on features of social recommender systems. *Artificial Intelligence Review*, 53(2), 2020.
- [142] David P. Shorten, Richard E. Spinney, and Joseph T. Lizier. Estimating transfer entropy in continuous time between neural spike trains or other event-based data. *PLOS Computational Biology*, 17(4):1–45, 04 2021.

- [143] Pranita Shrestha, Min Zhang, Yiqun Liu, and Shaoping Ma. Curiosity-inspired personalized recommendation. In *2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 33–40. IEEE, 2020.
- [144] Paul J. Silvia. *Exploring the Psychology of Interest*. Oxford University Press, New York, NY, USA, 2006.
- [145] Philipp Singer. *Modeling Aspects of Human Trails on the Web*. PhD thesis, Leibniz-Institute for Social Science, Knowledge Technologies Institute, Graz University of Technology, Graz, Austria, December 2014.
- [146] Philipp Singer. Modeling aspects of human trails on the web by philipp singer, with prateek jain as coordinator. *ACM SIGWEB Newsletter*, 1(Winter):3:1–3:2, jan 2016.
- [147] Alexandre M. Sousa, Jussara M. Almeida, and Flavio Figueiredo. Analyzing and modeling user curiosity in online content consumption: a lastfm case study. In *Proceed. of the 2019 IEEE/ACM Inter. Conf. on ASONAM*, ASONAM ’19, page 426–431, New York, NY, USA, 2020. ACM/IEEE.
- [148] Ajitesh Srivastava, Charalampos Chelmiss, and Viktor K. Prasanna. Influence in social networks: A unified model? In *Proceedings of the 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM ’14, page 451–454. IEEE Press, 2014.
- [149] Aji Subhanpurno and Ramachandra Kaimal. Document summarization using positive pointwise mutual information. *International Journal of Computer Science & Information Technology (IJCSIT)*, 2012.
- [150] Chenyu Sun, Hangwei Qian, and Chunyan Miao. From psychological curiosity to artificial curiosity: Curiosity-driven learning in artificial intelligence tasks. *arXiv preprint*, arXiv:2201.08300:1–35, 2022.
- [151] Jimeng Sun and Jie Tang. *A Survey of Models and Algorithms for Social Influence Analysis*, chapter 7, pages 177–214. Springer US, Boston, MA, 2011.
- [152] Jimeng Sun and Jie Tang. Models and algorithms for social influence analysis. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, WSDM ’13, pages 775–776, New York, NY, USA, 2013. Association for Computing Machinery (ACM).
- [153] Chenhao Tan, Jie Tang, Jimeng Sun, Quan Lin, and Fengjiao Wang. Social action tracking via noise tolerant time-varying factor graphs. In *Proceedings of the 16th*

- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '10, pages 1049–1058, New York, NY, USA, 2010. Association for Computing Machinery (ACM).
- [154] Jie Tang and Jimeng Sun. Computational models for social influence analysis. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14 Companion, pages 205–206, New York, NY, USA, 2014. Association for Computing Machinery (ACM).
- [155] Jie Tang, Jimeng Sun, Chi Wang, and Zi Yang. Social influence analysis in large-scale networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, pages 807–816, New York, NY, USA, 2009. Association for Computing Machinery (ACM).
- [156] Kanchana Thilakarathna, Aline Carneiro Viana, Aruna Seneviratne, and Henrik Petander. Design and analysis of an efficient friend-to-friend content dissemination system. *IEEE Transactions on Mobile Computing*, 16(3):702–715, 2017.
- [157] Nicholas M. Timme and Christopher Lapish. A tutorial for information theory in neuroscience. *eNeuro*, 5(3):1–40, 2018.
- [158] Natkamon Tovanich, Simone Centellegher, Nacéra Bennacer Seghouani, Joe Gladstone, Sandra Matz, and Bruno Lepri. Inferring psychological traits from spending categories and dynamic consumption patterns. *EPJ Data Science*, 2021.
- [159] K. E. Twomey and G. Westermann. Curiosity-based learning in infants: A neuro-computational approach. *Developmental Science*, 21(4):e12629, 2018.
- [160] Ashvanti Valji, Alisa Priemysheva, Carl J. Hodgetts, Alison G. Costigan, Greg D. Parker, Kim S. Graham, Andrew D. Lawrence, and Matthias J. Gruber. White matter pathways supporting individual differences in epistemic and perceptual curiosity. *bioRxiv*, 2019.
- [161] D. A. Vega-Oliveros and et al. The impact of social curiosity on information spreading on networks. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, ASONAM '17, page 459–466, NY, USA, 2017. IEEE/ACM.
- [162] Lasya Venneti and Aniket Alam. How curiosity can be modeled for a clickbait detector. *ArXiv*, abs/1806.04212, 2018.
- [163] Greg Ver Steeg and Aram Galstyan. Information transfer in social media. In *Proceedings of the 21st international conference on World Wide Web*, WWW' 12, pages 509–518, New York, NY, USA, 2012. Association for Computing Machinery (ACM).

- [164] Greg Ver Steeg and Aram Galstyan. Information-theoretic measures of influence based on content dynamics. In *Proceedings of the Sixth ACM international conference on Web Search and Data Mining*, WSDM '13, pages 3–12, New York, NY, USA, 2013. Association for Computing Machinery (ACM).
- [165] Carmen R. Wilson Van Voorhis and Betsy L. Morgan. Understanding power and rules of thumb for determining sample sizes. *Tutorials in Quantitative Methods for Psychology*, 3(2):43–50, sep 2007.
- [166] Chenxu Wang, Xiaohong Guan, Tao Qin, and Tao Yang. Modeling heterogeneous and correlated human dynamics of online activities with double pareto distributions. *Information Sciences*, 330:186–198, 2016.
- [167] Ningxia Wang, Li Chen, and Yonghua Yang. The impacts of item features and user characteristics on users' perceived serendipity of recommendations. In *Proceedings of the 28th ACM User Modeling, Adaptation and Personalization*, UMAP' 20, pages 266–274, New York, NY, USA, 2020. Association for Computing Machinery (ACM).
- [168] Zachary Wojtowicz and George Loewenstein. Curiosity and the economics of attention. *Current Opinion in Behavioral Sciences*, 35:135–140, 2020.
- [169] Chuhan Wu, Fangzhao Wu, Mingxiao An, Tao Qi, Jianqiang Huang, Yongfeng Huang, and Xing Xie. Neural news recommendation with heterogeneous user behavior. In *Proceedings of the 2019 EMNLP-IJCNLP*, pages 4874–4883, Hong Kong, China, nov 2019. Association for Computational Linguistics.
- [170] Q. Wu, S. Liu, and C. Miao. Modeling uncertainty driven curiosity for social recommendation. In *Proceedings of the International Conference on Web Intelligence*, WI '17, page 790–798, New York, NY, USA, 2017. Association for Computing Machinery (ACM).
- [171] Q. Wu and C. Miao. Curiosity: From psychology to computation. *ACM Computing Surveys*, 46(2):18:1–18:26, 2013.
- [172] Q. Wu, C. Miao, and B. An. Modeling curiosity for virtual learning companions. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '14, pages 1401–1402, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
- [173] Q. Wu, C. Miao, and Z. Shen. A curious learning companion in virtual learning environment. In *IEEE International Conference on Fuzzy Systems*, pages 1–8, Australia, 06 2012. IEEE.

- [174] Qiong Wu, Siyuan Liu, Chunyan Miao, Yuan Liu, and Cyril Leung. A social curiosity inspired recommendation model to improve precision, coverage and diversity. In *2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, pages 240–247. IEEE, 2016.
- [175] Qiong Wu and Chunyan Miao. Modeling curiosity-related emotions for virtual peer learners. *IEEE Computational Intelligence Magazine*, 8(2):50–62, 2013.
- [176] Xudong Wu, Luoyi Fu, Jingfan Meng, and Xinbing Wang. Maximizing influence diffusion over evolving social networks. In *Proceedings of the Fourth International Workshop on Social Sensing, SocialSense’19*, page 6–11, New York, NY, USA, 2019. Association for Computing Machinery (ACM).
- [177] W. M. Wundt. *Grundzude Physiologischen Psychologie*, volume 1 of *Duxbury Classic*. Wilhelm Engelmann, Leipzig, Germany, 1874.
- [178] K. Xu, J. Mo, Y. Cai, and H. Min. Enhancing recommender systems with a stimulus-evoked curiosity mechanism. *IEEE Trans. on Knowledge and Data Engineering*, 2019.
- [179] Ke Xu, Junwen Mo, Yi Cai, and Huaqing Min. Enhancing recommender systems with a stimulus-evoked curiosity mechanism. *IEEE Transactions on Knowledge and Data Engineering*, 33(6):2437–2451, 2021.
- [180] Xuenan Xu, Zeyu Xie, Mengyue Wu, and Kai Yu. Beyond the status quo: A contemporary survey of advances and challenges in audio captioning. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 32:95–112, oct 2023.
- [181] Andrew Yates, Rodrigo Nogueira, and Jimmy Lin. Pretrained transformers for text ranking: Bert and beyond. In *Proc. of the 14th ACM International Conference on Web Search and Data Mining, WSDM ’21*, pages 1154–1156, New York, NY, USA, 2021. Association for Computing Machinery.
- [182] Adel Jalal Yousif and Mohammed H. Al-Jammas. Exploring deep learning approaches for video captioning: A comprehensive review. *e-Prime - Advances in Electrical Engineering, Electronics and Energy*, 6:100372, 2023.
- [183] W. Youyou, M. Kosinski, and D. Stillwell. Computer-based personality judgments are more accurate than those made by humans. *National Academy of Sciences*, 112(4):1036–1040, 2015.
- [184] Chenhui Zhang, Sida Gao, Jie Tang, Tracy Xiao Liu, Zhanpeng Fang, and Xu Cheng. Learning triadic influence in large social networks. In *IEEE/ACM ASONAM*, pages 1380–1381, Davis, California, USA, 2016. IEEE/ACM.

-
- [185] Lejun Zhang, Tong Wang, Zilong Jin, Nan Su, Chunhui Zhao, and Yongjun He. The research on social networks public opinion propagation influence models and its controllability. *China Communications*, 15(7):98–110, 2018.
 - [186] P. Zhao and D. L. Lee. How much novelty is relevant?: It depends on your curiosity. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, page 315–324, New York, NY, USA, 2016. Association for Computing Machinery (ACM).
 - [187] Yuqing Zhu, Jing Tang, and Xueyan Tang. Pricing influential nodes in online social networks. *Proceedings of the VLDB Endowment*, 13(10):1614–1627, June 2020.
 - [188] Perry Zurn and Danielle S. Bassett. On curiosity: A fundamental aspect of personality, a practice of network growth. *Personality Neuroscience*, 1:e13, 2018.

Appendix A

Axioms for Traditional Collative Variables

We introduce axioms that deal with the basis for understanding the behavioral patterns of the values of our metrics for the traditional collative variables of curiosity. Let there be a random variable X and a probability $P_t^\rightarrow$ such that $P_t^\rightarrow(X = x)$ is the probability of a outcome x from random variable X .

Definition 1 (Axiom of monotonicity decreasing for surprisal (MONO-SURP)). *A metric based on the surprisal, i.e., Shannon information content, using $P_t^\rightarrow(X = x)$ is monotonically decreasing for all $P_t^\rightarrow(X = x) \in [0, 1]$.*

The axiom MONO-SURP can be interpreted as follows. Let the following probabilities be $P_t^\rightarrow(X = x_1)$ and $P_t^\rightarrow(X = x_2)$, where x_1 and x_2 are two different outcomes of the same random variable X , for which a metric such as $-\log_2 P_t^\rightarrow(X = x_1)$ and $-\log_2 P_t^\rightarrow(X = x_2)$ is calculated, then $\forall P_t^\rightarrow(X = x_1), P_t^\rightarrow(X = x_2) \in [0, 1]$

$$\text{if } P_t^\rightarrow(X = x_1) > P_t^\rightarrow(X = x_2) \Rightarrow -\log_2 P_t^\rightarrow(X = x_1) < -\log_2 P_t^\rightarrow(X = x_2).$$

Definition 2 (Axiom of maximization and minimization for entropy (MAX-MIN)). *A metric based on Shannon entropy that uses $P_t^\rightarrow(X = x)$ as the class probability of n classes is maximized if the distribution of probabilities is uniform, i.e., $P_t^\rightarrow(X = x) = 1/n$. Otherwise, if there is $P_t^\rightarrow(X = x) = 1$ for one of classes, the metric is minimized and equals 0.*

Intuitively, the MAX-MIN axiom corresponds to the properties of Shannon entropy [92, 32]. Let $H(X)$ be a function that represents the Shannon entropy¹ for n classes (e.g. musical genres), then let $0 \leq H(X) \leq \log_2 n$. Therefore, $H(X) \geq 0$ with equality if and only if there is a class with $P_t^\rightarrow(X = x) = 1$,

$$\forall n \in \mathbb{N}^+ \wedge \forall P_t^\rightarrow(X = x) \in]0, 1], \text{ if } \exists P_t^\rightarrow(X = x) = 1 \Rightarrow H(X) = 0,$$

and $H(X) \leq \log_2 n$ with equality if for all classes $P_t^\rightarrow(X = x) = 1/n$, that is

$$\forall n \in \mathbb{N}^+ \wedge \forall P_t^\rightarrow(X = x) \in]0, 1[, \text{ if } P_t^\rightarrow(X = x) = 1/n \Rightarrow H(X) = \log_2 n.$$

¹Shannon entropy is the average Shannon information content.

To show the connection of the axiom **MAX-MIN** with the metric based on Shannon entropy, we use the Kullback-Leibler divergence² (D_{KL}) [92] between two probability distributions to test whether the probability distributions of a class (i.e., musical genres) are uniform. So, if $D_{KL} = 0$, then the probabilities are uniformly distributed.

For metrics based on Surprisal, the axiom **MONO-SURP** describes a behavioral pattern, namely that the metrics decrease monotonically depending on their probabilities. This means that the more often the user accesses a certain content in a short period of time (i.e. within the window of interaction), the less stimulating this content is for him in terms of curiosity in the long term. Conversely, the less often the user accesses a content (in a short period of time), the greater their curiosity about this content item (in the long term).

The **MAX-MIN** axiom in turn shows the upper and lower bounds for metrics based on Shannon entropy. This means that the more uniformly the user's access to the different classes of content items is distributed over a short period of time (i.e., within the window of interaction), the more stimulating these classes of content item are for him in terms of curiosity in the long term. This means that users tend to consume items in these classes in the near future because they are more interesting to them.

In short, we emphasize that the metrics defined for the traditional collative variables are always non-negative and should be interpreted as: *the greater the value the higher the stimulus to the curiosity of the user*. In the next sections we describe some examples of the relationship between metrics and axioms.

A.1 The LastFM Case Study: examples

The axiom **MONO-SURP** helps us to better understand the metric based on the surprisal and shows us the behavior of the possible values that each metric can capture in the case study of LastFM. For this purpose, we chose a user with the largest number of listening events, namely 173,508, and arbitrarily selected all the listening events where $|C_t^+| = 5$. Figure A.1 shows the behavior of the values in relation to the **MONO-SURP** axiom for the metrics *songNovelty*, *artistNovelty* and *genreNovelty* of this user. Figure A.1a shows the minimum and maximum values for *songNovelty*, which are 3.67 bits and 8.13 bits. For *artistNovelty* in Figure A.1b, the minimum is 0 bits (no information) and the maximum is 8.13 bits. Finally, for *genreNovelty* a minimum of 0.5 bits and a maximum of 6.74 bits is displayed in Figure A.1c. It can be seen in these figures that the values

²Kullback-Leibler divergence or relative entropy between two probability distributions $P(x)$ and $Q(x)$ is $D_{KL}(P||Q) = \sum_x P(x) \log_2 P(x)/Q(x)$.




Figure A.1: Relation of *songNovelty*, *artistNovelty* and *genreNovelty* and their respective probabilities for the axiom MONO-SURP.

of the metrics decrease with increasing probability, which also shows the agreement with the axiom MONO-SURP.

Again, we use the same user from the example in Figure A.1 to show that the latter three surprisal-based metrics are also consistent with the axiom MONO-SURP. Figure A.2 thus shows the correlation of *instComplex*, *overallComplex* and *genreConflict* with this axiom. We see that this figure shows the same monotonic behavior of the first three metrics also for the last three metrics, i.e. the lower the probability (x -axis), the higher the surprisal (y -axis) for the two metrics of complexity (instantaneous and overall) and for the conflict for all musical genres in the window of interaction.




Figure A.2: Relation of *instComplex*, *overallComplex* and *genreConflict* and their respective probabilities for the axiom MONO-SURP.

Finally, Figure A.3 thus shows an example for the selected user with $|C_a| = 5$ (also used in Figure A.1), where the y -axis represents *genreUncertainty* and the x -axis D_{KL} . In this figure, the minimum value for *genreUncertainty* is 0.78 bits and the probabilities for each class is $\{0.01, 0.01, 0.13, 0.01, 0.84\}$. So, we can see that these probabilities are far from the uniform distribution where each probability is $1/5$, because that, it has $D_{KL} = 1.07$, the largest value. On the other hand, the Shannon entropy for *genreUncertainty* reaches a maximum of 2.32 bits at $D_{KL} = 0$. It can be seen that *genreUncertainty* decreases monotonically in relation to D_{KL} . This means that the closer D_{KL} is to 0, the




Figure A.3: Relation of $genreUncertainty$ and D_{KL} value for the axiom MAX-MIN.

greater $genreUncertainty$. In other words: If D_{KL} tends towards 0, the distribution of the probabilities of the musical genres comes closest to a uniform distribution.

A.2 The WhatsApp Case Study: examples




Figure A.4: Relation of $userNovelty$ and $catNovelty$ and their respective probabilities for the axiom MONO-SURP.

In this section, we use the axioms MONO-SURP and MAX-MIN to show how the traditional metrics, i.e., the collative variables, also behave in the case study of WhatsApp. We therefore select the most active user with 18,879 shared messages from the WhatsApp group with the largest number of users (1,798) and shared messages (415,013). The Figures A.4a and A.4b show how the metrics $userNovelty$ and $catNovelty$ decrease monotonically with respect to the increasing values of their probabilities, and this is consistent with the axiom MONO-SURP. In addition, the maximum and minimum values are highlighted in these figures. In Figure A.4a, for example, the maximum value is 6.15 bits for $P_{t,g}^{\rightarrow}(U=u) = 20\%$ (the lowest probability) and for $P_{t,g}^{\rightarrow}(U=u) = 76\%$ (the highest




Figure A.5: Relation of *userUncertainty* and *catUncertainty* and their respective D_{KL} values for the axiom MAX-MIN.

probability) the minimum value is 0.39 bits. Finally, Figure A.4b shows 1.21 bits for $\bar{P}_{t,g}^{\rightarrow}(\mathcal{C}_m) = 43\%$, which is the maximum value, and the minimum value is 0 bits – no surprisal – for 100% of the probability.

The Figure A.5a (A.5b) show the relationship between the uncertainty of the user *userUncertainty* (the category of messages *catUncertainty*) and the axiom MAX-MIN. This axiom states that Shannon entropy is maximized when the probability distribution of a class (user or category of messages) is uniform. For this purpose, the Kullback-Leibler divergence (D_{KL}) is used. In this way, the figures show D_{KL} on the x -axis and the values of these metrics on the y -axis. Figure A.5a shows, for example, all values of *userUncertainty* for all windows of interactions with exactly 5 users of the active user of the group with the largest number of messages. This figure shows that the maximum value of *userUncertainty* = 2.3 bits is reached when $D_{KL} = 0.02$, it means that the distribution of users in the window of interaction is almost uniform. In contrast, the minimum value is reached when $D_{KL} = 0.81$, that is, *userUncertainty* = 1.14 bits. The same behavior can be observed for *catUncertainty*, in this case, we arbitrarily select all windows of interaction which has at most 4 types of message categories. The maximum value is 1.79 for $D_{KL} = 0.14$, which brings the probability distribution closer to the uniform distribution. The distribution that is further away from the uniform distribution has $D_{KL} = 1.37$ and reaches the value of 0.13 for *catUncertainty*.

The Figures A.6a and A.6b show the relationship between the conflict metrics and the axiom MONO-SURP. This means that the values of *userConflict* and *catConflict* (y -axis) decrease monotonically as the values of the probabilities increase (x -axis). We then use the same example of the most active user from the group with the most shared messages, we get all the windows of interaction in which the number of users is 83. So, the maximum value for *userConflict* is 4.06 bits for a probability of 6% (the lowest value), while the minimum value is 2 bits for probability of 25%. Similarly, Figure A.6b shows the behavior of *catConflict* according to the axiom MONO-SURP and that the windows




Figure A.6: Relation of *userConflict*, *catConflict* and *catComplex* and their respective probabilities for the axiom MONO-SURP.

of interaction have exactly 83 shared messages. In this figure, we see that the lowest probability is 20% and reaches 2.34 bits, while the highest probability is 100% and reaches 0 bits, i.e. no information or no surprisal.

Figure A.6c shows, for example, the relationship between the axiom MONO-SURP and the complexity, the last metric. In this figure, the maximum value for *catComplex* is 2.32 bits at the lowest probability of 20% and the minimum value of 0 bits at the highest probability of 100%, which is no surprisal (i.e., no information). This shows that *catComplex* is monotonically decreasing with respect to its probabilities.

Appendix B

Axioms for Social Curiosity Variables

We present axioms that form the basis for understanding behavioral patterns of the values of social curiosity metrics. Let random variables X and Y , of which $P(X=x)$ is the probability of the random variable $X = x$ and $P(X=x|Y=y)$ is the conditional probability of $X = x$ given $Y = y$, then Pointwise Mutual Information (PMI) [19] is defined as

$$PMI = \log_2 \frac{P(X = x|Y = y)}{P(X = x)}.$$

Definition 3 (Axiom of increasing monotonicity for PMI (MONO-PMI)). *A metric based on PMI is monotonically increasing with respect to its ratio $P(X = x|Y = y)/P(X = x)$ for all $P(X = x|Y = y) \in [0, 1]$ and $P(X = x) \in [0, 1]$.*

We can understand the axiom MONO-PMI as follows. Let $f(z)$ be a function denoting the PMI, and let z denote the ratio underlying the PMI, i.e., $P(X = x|Y = y)/P(X = x)$, then $\forall z, w \in [0, 1]$, if $z \leq w \Rightarrow f(z) \leq f(w)$.

Yet, the minimum and maximum values can be defined as follows. PMI is maximum when X and Y have the greatest co-occurrence, i.e., $P(X = x|Y = y) = 1$ or $P(Y = y|X = x) = 1$ [20], which gives the following upper bound

$$PMI \leq \min(-\log_2 P(X = x), -\log_2 P(Y = y)).$$

The lower bound [20] is given as follows

- (i) if $P(X=x|Y=y)/P(X=x) = 1$, then X and Y are independent and PMI is 0;
- (ii) if $P(X=x|Y=y) = 0$, then $P(X=x, Y=y) = 0$ and $P(X=x) > 0$, so PMI is 0;
- (iii) if $P(X=x|Y=y) < P(X=x)$, then PMI is negative, so we can set the metric based on PMI to 0.

This results in the following lower and upper bounds

$$0 \leq PMI \leq \min(-\log_2 P(X = x), -\log_2 P(Y = y)).$$

For the following axiom, the mutual information [92, 32] between two random variables X and Y is given as

$$MI(X, Y) = \sum_{y \in Y} \sum_{x \in X} P(X=x, Y=y) \log_2 \frac{P(X=x|Y=y)}{P(X=x)},$$

and let $H(X|Y)$ be the conditional entropy defined as follows

$$H(X|Y) = \sum_{x \in X} \sum_{y \in Y} P(X=x, Y=y) \log_2 \frac{P(X=x, Y=y)}{P(X=x)}.$$

Definition 4 (Axiom of lower and upper bounds for the mutual information (BOUND-MI)). *A metric based on mutual information has the following lower and upper bounds $0 \leq MI(X, Y) \leq H(X)$.*

The upper bounds are defined as follows $MI(X, Y) \leq H(X)$, if $MI(X, Y) = H(X) - H(X|Y)$, then $MI(X, Y) \leq H(X)$, with equality if and only if X is completely determined by Y , i.e., $H(X|Y) = 0$ [92, 32].

According to the properties of mutual information [32], the lower bound is given as $MI(X, Y) \geq 0$, with equality if and only if X and Y are independent, i.e., $P(X = x, Y = y) = P(X = x)P(Y = y)$.

Finally, the lower and upper bounds for MI are defined as follows

$$0 \leq MI(X, Y) \leq H(X).$$

For our last axiom definition, let $MI_{pt}(X, Y = y)$ be the partial mutual information for a given $y \in Y$, defined as

$$MI_{pt}(X, Y = y) = \sum_{x \in X} P(X=x, Y=y) \log_2 \frac{P(X=x|Y=y)}{P(X=x)}.$$

Definition 5 (Axiom of lower and upper bounds for partial mutual information (BND-PTMI)). *A metric based on partial mutual information has the following lower and upper bounds $0 \leq MI_{pt}(X, Y = y) < MI(X, Y)$.*

If X and Y are dependent, there is an upper bound for MI_{pt} , which is defined as $MI_{pt}(X, Y = y) < MI(X, Y)$. In contrast, the lower bound for MI_{pt} is given as follows

- (i) if X and Y are independent, i.e., $P(X = x, Y = y) = P(X = x)P(Y = y)$, for all $x \in X$ and for a given $y \in Y$, then $MI_{pt}(X, Y = y) = 0$;
- (ii) if $P(X = x|Y = y) < P(X = x)$, then MI_{pt} is negative, so we set the metric based on MI_{pt} to 0.

The lower and upper bounds for the partial mutual information are therefore as follows

$$0 \leq MI_{pt}(X, Y = y) < MI(X, Y).$$

In the next section, we describe some examples of the relationship between the social curiosity metrics and their respective axioms from the WhatsApp case study.





Figure B.1: The axiom MONO-PMI: relationship between the metrics of direct social influence and the ratio $P_{t,g}^{\leftarrow}(U=u|S=s)/P_{t,g}^{\leftarrow}(U=u)$.

B.1 The WhatsApp Case Study: examples

The axiom MONO-PMI shows us how the possible values of the metrics for direct social influence behave. We therefore select the most active user with 18,879 shared messages from the WhatsApp group with the largest number of users (1,798) and shared messages (415,013). Figure B.1a shows the relationship between $maxDirInf$ (blue line) and the ratio $P_{t,g}^{\leftarrow}(U=u|S=s)/P_{t,g}^{\leftarrow}(U=u)$ (x-axis) with the MONO-PMI axiom, which states that PMI increases monotonically with respect to its ratio. The more the ratio increases, the more $maxDirInf$ increases. The minimum and maximum values of $maxDirInf$ are highlighted in the figure. In addition, for each value of the ratio on the x-axis, this figure shows the corresponding value for $avgDirInf$ (red circles). It can be seen that the red circles are below the blue line of $maxDirInf$, which serves as the upper limit for this value. To confirm this, Figure B.1b shows the distribution of these two metrics and it can be seen that the distribution of $maxDirInf$ is larger than the distribution of $avgDirInf$ for 90% of the values and is the same for the rest (10%).

In the following, we show how the metrics of indirect social influence are correlated. To do this, we return to the user who is the most active group from the example in Figure B.1. Figure B.2a shows a scatterplot for each pair (x, y) of the indirect social influence metric, namely $maxIndInf$ versus $avgIndInf$. In this figure, the minimum value achieved by this user is highlighted with the value (0,0) from both metrics. In addition, the largest values for the two axes are highlighted, with 0.064 for the $avgIndInf$ (x-axis) and 0.084 for the $maxIndInf$ (y-axis). Yet, it can be seen that the values for $maxIndInf$ are larger than for $avgIndInf$ and the blue solid line shows the threshold at which the values of $avgIndInf$ are limited. Finally, the Figure B.2b shows the distributions for these two metrics and for more than 90% of the values the $maxIndInf$ is greater than



(a) $maxIndInf$ versus $avgIndInf$. (b) Distributions of metrics.

Figure B.2: Relationship between the metrics of indirect social influence.

the $avgIndInf$.

To establish a relationship between the metric of individual indirect social influence and the group-level metrics, the Figure B.3 shows a scatterplot in which each circle represents a pair of metrics as follows: (a) the *green circles* denotes the $avgIndInf$ versus $maxIndInf$; (b) the *blue circles* represents the $avgIndInf$ versus $groupMutInf$; (c) the *yellow circles* denotes the $avgIndInf$ versus $H_{t,g}^{\leftarrow}(U)$ (entropy of curious users); and (d) the *red circles* represents the $avgIndInf$ versus the maximum entropy of curious users. We emphasize that this figure also shows how the axiom BOUND-MI and BND-PTMI are related, because $groupMutInf$ is an upper bound for $avgIndInf$ and $maxIndInf$ and $H_{t,g}^{\leftarrow}(U)$ is an upper bound for $groupMutInf$. In Figure B.3 we see that the set of *green circles* ($maxIndInf$) is below the *blue circles* ($groupMutInf$) and the set of *yellow circles* is above the *blue circles*. This pattern applies to all users in the WhatsApp dataset and shows the boundaries of the values for the individual metrics and the group-level metrics. At the top of this figure, we also plot the *red circles* as the maximum entropy of *curious users* that would be achieved if the distribution of *curious users* (i.e., influenced users) were uniformly distributed, and it shows the upper bound for all metrics.




Figure B.3: Relationship between the metrics and the axiom BOUND-MI and BND-PTMI.