Delft University of Technology

Faculty of Electrical Engineering, Mathematics and Computer Science
Computer Science
Web Information Systems

Documentation

Identifying scientific memes and their interactions in online communities

Joost Verdoorn

Committee

Prof.dr.ir. Geert-Jan Houben (Chair)

Dr. Nava Tintarev (Daily advisor)

Dr. Huijuan Wang

Dr. Christoph Lofi (Daily co-advisor)

May 14, 2018

Joost Verdoorn
Identifying scientific memes and their interactions in online communities
Documentation, May 14, 2018
Committee: Prof.dr.ir. Geert-Jan Houben, Dr. Nava Tintarev,
Dr. Huijuan Wang, Dr. Christoph Lofi

Delft University of Technology
Web Information Systems
Computer Science
Faculty of Electrical Engineering, Mathematics and Computer Science
Mekelweg 2
2628 CD and Delft

Abstract

Memes are theorized to be the building blocks of culture. Due to a lack of empirical validation, however, the theory of memes — *memetics* — remains in its infancy. We argue that one of the missing components for such empirical validation is a method for the large-scale identification of memes.

In this thesis, we develop a method for the identification of *scientific memes* — ngrams of length 1 through 4, denoting scientific concepts — propagating within online communities. With data extracted from science-oriented correspondence extracted from five communities on the online discussion platform Reddit, and five communities on the online question and answer platform StackExchange, we perform a large-scale automated evaluation in which we find that memes identified in these communities correspond to the titles of Wikipedia articles; and a small-scale human evaluation in which we find that the identified memes represent relevant concepts to the community's scientific field.

Furthermore, we introduce a slight adaptation of this method to elucidate one of memetics' predictions: the occurrence of *interactions* between memes, where the occurrence of one meme has a positive or negative influence on the propagation of another meme. To evaluate this method for the identification of meme interactions, we construct *meme interaction networks*, in which we find that the most central memes correspond to the most relevant scientific concepts.

We find that our methods are able to extract key concepts within online communities, identifying thousands of relevant concepts from millions of candidate ngrams. Thus, our method may contribute to contemporary text mining research, and could be used in place of, or in conjunction with current approaches, such as TF-IDF or LDA.

Acknowledgement

Firstly, I'd like to thank my primary supervisor Nava Tintarev for her guidance throughout my thesis. She's an inspiring academic and a wonderful person, never shying away from giving fair criticism when things go wrong, but also giving praise when things go right. She has taught me a lot about the academic process and being critical of my own work and that of others. Nava, I am very fortunate to have had you as my thesis supervisor.

Many thanks also to my secondary supervisor Christoph Lofi. Often providing a contrasting view to Nava, they as a team allowed me to see the spectrum of potential solutions beyond my own intuitions. Christoph, thank you for taking the time listening to my ramblings, but never accepting them at face value — your practical way of thinking brought some much needed balance to the early work of the thesis.

Further, I'd like to thank my dearest Fenna van der Wal for being patient and supportive throughout my studies, and particularly during my thesis project. She kept me sane, being able to pull me back to reality when I was stuck on one of the many problems encountered, but also providing me with the encouragement to get back to work when things were tough.

I'm grateful for Steffan Sluis and our early discussions on information, knowledge, and memes. I know I've annoyed you to death with my constant ramblings about memes, but now that it's finished I know you can start to forgive me.

Jac van Egmond for listening, offering advice, and providing some much needed distraction and companionship. Always keen to know how it was going, but also never afraid to explore and discuss other topics close to our hearts.

And finally, my dear friends and colleagues at FeedbackFruits, the educational technology company which provided me with food and shelter, laughter and discipline. I'm very much looking forward to finally joining you guys in a full time capacity. Let's change the world together!

Contents

1 Introduction

> *Memes: viral concepts, infections of conscious thought. Some flare and die like mayflies. Others last a thousand years or more, tricking billions into the endless propagation of parasitic half-truths.*
>
> — **Peter Watts**
> Author of 'Maelstrom'

1.1 Motivation

In this thesis, we take it upon us to identify patterns in behavior — *memes* — which propagate through online, science-oriented communities. Memes herein take the form of distinct textual patterns — ngrams of length 1 through 4, for example *recursion*, *lambda calculus*, and *np-complete* — which we trace as they move from user to user. Furthermore, we look at the interaction that these memes have with one another as they propagate through the communities, with some memes propagating better under each other's influence. The identified memes and their interactions allow us to extract key concepts from online communities.

Memes are theorized to be the building blocks of culture. The study of memes — *memetics* — describes how these memes propagate from person to person, spread through a population, and evolve over time. Memetics combines aspects from the disparate fields of social contagion, cultural evolution, and information diffusion into a coherent framework, providing an explanation for the propagation, spread, and evolution of culture. However, the theory of memetics remains in its infancy, as empirical validation is still largely lacking, and many of its predictions have yet to be put to the test. One such a prediction is that memes engage in complex interactions with one another, leading to the formation of clusters of culture, such as scientific fields, religions, and political ideologies.

We identify some of the key missing components to the validation of memetics as a historical lack of data, and the absence of a general method to identify memes and their interactions on a population scale. With more and more data on human correspondence being available for a myriad of online communities, we are inspired to address the following research questions:

RQ1: Can we identify memes propagating through online communities?
RQ2: Can we identify the interactions between these memes?

To further show the relevance of memetics to contemporary scientific endeavor, and contribute to the discipline of information extraction, we will also address the following research question:

RQ3: Can we leverage memetic theory to extract key concepts from online communities?

It is our hope that through addressing these research questions, we will provide a stepping stone for further research into memetics. We argue that the methodologies we develop here can also be applied to identify memes in other types of human correspondence data, as long as we can determine two key properties: what behavior was an individual exposed to, and what behavior did they express themselves. These key pieces of information allow us to identify memes and their interactions by quantifying the effect that exposure to some behavior has on the subsequent expression of the same or other behavior, respectively. The proposed methods can be applied to identify memes on a population scale, with our evaluations identifying memes in communities consisting of hundreds of thousands of unique users and posts, selecting a few thousand memes from tens of millions of unique ngrams.

1.2 Results

In order to address the research questions stated above, we introduce the Memeseeker method for the large-scale extraction of textual memes from correspondence data of online communities. This method combines insights from previous works in memetics, social contagion, and epidemiology into a three-step meme identification process.

Through an automated evaluation of the Memeseeker method, applied to correspondence data of online platforms Reddit and StackExchange, we find that the method is able to extract *scientific memes* — here, representing titles of Wikipedia articles — with a reasonably high precision, and ranking them reasonably well. Furthermore, through a more fine-grained human evaluation, we find that the ngrams our method identifies as memes correspond to relevant topics and concepts within the particular scientific fields the communities from which they are extracted represent.

We further suggest a slight adaptation to the Memeseeker method in order to identify not only memes, but also the interactions thereof. To overcome the inherent difficulty in evaluating this method due to the novelty of the task, we introduce the notion of *meme interaction networks*, and calculate the *centrality* of memes within such a network. Our evaluations find that meme centrality within meme interaction networks is a better predictor for human-judged relevance of terms to their scientific field, compared to meme centrality in trivially-constructed meme co-occurrence networks.

Lastly, through both our automated and human evaluations, we find that our method is able to extract key concepts to particular communities from the correspondence data produced by such communities, with these key concepts being represented by the memes that propagate through such communities.

1.3 Contribution

Foremost, our Memeseeker method demonstrates that it is possible to identify *scientific memes* that propagate through science-oriented online communities, on a web scale. Future research should evaluate if this is the case for other types of memes, but we see no theoretical hurdles.

Secondly, we find support that a slight adaptation of the Memeseeker allows us to not only identify memes, but also their interactions. These interactions further allow us to construct interaction networks for the identified memes, which we argue could potentially be a very useful tool for the study of culture. Such networks, which are analogues to interaction networks in other domains, such as gene and protein interaction networks, and ecological networks, should allow for future investigation of one of memetics' predictions — the formation and evolution of co-adapted meme complexes — and, as we will theorize, *meme pathways*.

Furthermore, we find that our Memeseeker method may provide an alternative to keyword extraction approaches such as TF-IDF, extracting important concepts from online correspondence on a web scale, and that it may contribute to social network analysis through laying bare the influence pathways over which imitation takes place.

1.4 Thesis Structure

Chapter 2 We discuss work related to the topics discussed in this thesis. In particular, we will discuss theoretical work related to memetics, and the identification of memes and their interactions. Furthermore, we will shortly discuss the most common methods for keyword and topic extraction.

Chapter 3 We build upon the related work laid out in Chapter 2 to arrive at the Memeseeker method for the identification of memes. We further demonstrate a slight adaptation of this method which allows us to not only identify memes, but also the interactions in which they engage.

Chapter 4 We evaluate our Memeseeker method for the identification of memes on five datasets extracted from the online discussion platform Reddit. We first perform a coarse, large-scale offline evaluation, and then move on to perform a finer, but small-scale evaluation using human participants.

Chapter 5 We evaluate the adapation of the Memeseeker method to identify the interactions in which memes engage. For five datasets extracted from Reddit, we construct a meme interaction network and a co-occurrence network. We then evaluate whether centrality in meme interaction networks is a better predictor for relevance than centrality in co-occurrence networks.

Chapter 6 We return to the research questions as stated at the beginning of this Chapter, and will answer them with respect to the results obtained through the evaluations in Chapters 4 and 5.

Related Work

2

> *People don't have ideas. Ideas have people.*
>
> — **Carl Jung**
> Psychiatrist

2.1 Introduction

In this Chapter, we first introduce the theoretical work on memes in Section 2.2; we will then discuss work related to the identification of memes in Section 2.3; after which, we will move to discuss works related to meme interactions in Section 2.4; in other fields, similar interactions have led to the construction of interaction networks, which we will discuss in Section 2.5; and finally, we will discuss some common methods for the extraction of concepts and topics from text corpora in Section 2.6.

2.2 Memes

In 1976, Richard Dawkins released his seminal work The Selfish Gene [Daw76]. Through this work, he argues that biological evolution does not take place on the level of the *organism* — survival of the fittest individuals — but on the level of the *gene* — survival of the fittest genes. Furthermore, he extrapolates this notion to arrive at the concept of *Universal Darwinism*: evolution can and will happen wherever there is some entity that adheres to three fundamental properties: it has *heredity* of traits, has *variance* of these traits, and there is *selection* of these traits. He calls this entity a *replicator*, which he posits as the *unit of natural selection*.

Dawkins further notes that the gene is not the only replicator we see in the world around us. Culture — he argues — is also subject to evolution, and this evolution is also driven by a replicator, which he calls the *meme*. Memes — originally short for *mimeme*, or that which is imitated — are patterns in behavior that are passed on by imitation. According to Dawkins, memes are active replicators: there is *inheritance* of behavioral patterns as memes replicate from person to person through imitation; this replication process may be imperfect, causing *variance* of those behavioral patterns; and there is *selection* of behavioral patterns, as not every behavioral pattern is imitated. These properties — inheritance, variance, and selection — cause memes to evolve over time, as new behavior is invented, imitated and adapted.

One of the critiques the theory of memetics has consistently received over the years is that it cannot define the exact unit of a meme [Bla00; Shi13]. This means that, although we do recognize there is *some* pattern in behavior that is being spread through imitation, it is very hard to delimit where one meme starts and the other stops, or whether we are actually looking at a set of memes instead of a single meme. Indeed, arguably the largest hurdle for memetics is the current inability to identify memes in the wild, such that they can be empirically studied, inspire new theoretical developments, and memetics' predictions can be tested. In the next Section, we will detail work which has been done toward such a method.

2.3 Identification of memes

Recently, many publications in a number of disciplines have used the notion of memes. However, because a general identification method for memes is still lacking, many of these works have resorted to using hand-picked operational proxies for memes: behavioral patterns that the researchers qualify as being memes, such as hashtags on Twitter [TR12; TR15; Wen+14], topical clusters [CC13], or even manually classified tweets [Car+15]. The usage of these hand-picked proxies unfortunately limits the applicability of this research to domains where such proxies are unavailable.

There have been a number of proposed methods for the large-scale identification of memes. Notably, Leskovec et al. [Les+09] and Suen et al. [Sue+13] use overlapping textual phrases found in blog posts and news articles. Their research provides insights into the dynamics of the news cycle, as key news phrases — textual memes — are continuously created, replicated, and adapted as they are picked up by news sites and blogs. Although their method allows for the reconstruction — and thus identification — of textual memes, their method is limited to relatively coarse textual fragments, as reconstruction becomes less reliable for shorter fragments.

Christakis and Fowler [CF07; CF08; CF13] investigate whether deleterious behavioral traits, such as obesity, smoking, and substance abuse, spread contagiously through social networks. The authors quantify the strength of contagion of such behavioral traits through comparison of the probability of an individual expressing some trait when they either are or are not exposed to an acquaintance expressing that trait. Through their work, the authors find that a number of behavioral traits indeed spread socially, or when viewed through the lens of memetics, they find that some behavioral traits behave like memes. However, coming from an epidemiological background, the authors only look at some pre-determined set of behavioral traits, for which pre-existing long-running longitudal medical studies can provide data.

In a similar fashion, Kuhn et al. [Kuh+14] investigate the propagation of simple words and phrases — which they call *scientific memes*, represented as ngrams of length 1 through 4 — through citation networks of scientific literature, by comparing the probability of a publication containing some ngram when it either does or does not cite a work containing that ngram. However, instead of performing such analysis for some pre-selected set of ngrams — as Christakis and Fowler did for some pre-selected set of behavioral traits — the authors perform it for every ngram present in the corpus, which allows them to find those ngrams that are most contagious, and thus, behave most like memes as they propagate through the citation network. This method, being grounded in the replication mechanism of memes, allows the authors to extract memes from corpora with a much higher granularity than previous methods, like the ones proposed by Leskovec et al. or Suen et al., although the reliance on citation networks does limit the applicability to other types of corpora.

Krawczyk and Kułakowski [KK16] build on the work of Kuhn et al., and find empirical support for the the fact that not only scientific literature itself, but also its authors can be regarded as *vehicles* for memes. This result supports the distinction memetics makes between the people who spread memes (*egos*) and the *artefacts* through which those memes spread, but regarding both as vehicles for memes [Bla00].

More recently, Beck-Fernandez et al. [BF+17] propose a method for the identification of memes in online fora. Their method is based upon the identification of relationships between pairs of words, some of which, the authors argue, can be considered memes. However, the method is not grounded in memetic theory, and we are not convinced such word relations correspond to the memes described by memetic theory. Furthermore, as the authors use semantic and syntactic knowledge (using e.g., WordNet), this hinders the applicability of their approach to specialized or non-English corpora for which concepts and relations are unknown beforehand.

2.4 Identification of meme interactions

Memetics predicts that memes typically do not exist in isolation, but form complex interactions with one another [HC09]. Since the inception of memes by Dawkins [Daw76], memes have been theorized to form co-adapted meme complexes — *meme-*

plexes — resembling clusters of culture such as scientific fields, political idealogies, and religions [Bla00]. The process by which such complexes are thought to arise — *meme interactions* — are thought to be either positive, where the existence of a meme positively impacts the expression of another meme; negative, when this impact is negative; or neutral, when there is no significant interaction [Bes97]. Groups of memes which have a mutually positive interaction are theorized to form memeplexes [Bla00].

Best [Bes97] was the first to propose a quantification of the interactions between memes. Through measurement of meme occurrence frequencies on message boards over time, they determined the cross-correlation between such occurrence frequencies for a number of meme pairs. A meme for which the occurrence over time showed a high cross-correlation with the occurrence of another meme was said to positively influence that meme. Similarly, negative influences were also detected. This result was inspiring for memetics, finding support for its predictions that memes engage in interactions with one another. Nonetheless, the method does have some downsides: by looking only at raw meme occurrence in fixed time intervals, the granularity of the interactions that can be identified through this method is relatively coarse, failing to pick up interactions which may play out over longer time intervals. Besides, noisy data may mask the finer correlations that occur, further limiting the applicability of this method.

Chavalarias and Cointet [CC13] use simple meme co-occurrence in scientific literature to get a sense of the interactions between such memes. This allows them to cluster memes which often co-occur into groups of memes — *memeplexes*, in this case scientific disciplines — which they then trace over time, as these memeplexes emerge, branch, merge, and disappear. This allows them to trace the evolution of scientific fields over time. Though word co-occurrence is often used in text mining tasks to get a sense of the relations between words, it may not be able to identify the finer interactions memes engage in. Foremost, co-occurrence is a *symmetric* measure, whereas memetics predicts meme interactions can be asymmetric. Furthermore, co-occurrence is limited to singular pieces of text (documents), whereas meme interactions may play out across a corpus. Nonetheless, the simplicity of this method may make it a viable way to determine some of the coarse interactions in which memes engage.

2.5 Interaction networks

Interaction networks are networks which model the interactions within complex systems. Although interaction networks have thus far not seen any use in memetics, they are common in other fields, such as in proteomics (e.g. protein interaction network), genomics (gene interaction networks), and ecology (ecological networks). In those disciplines, interaction networks allow for the fine-grained analysis of interaction patterns, and the emergent properties these interactions result in. Similar analyses may be applied to meme interaction networks, and we will discuss some below.

In ecological networks, *centrality analysis* have been used to identify *keystone species*, important species to the local ecology [Jor+06]. In protein interaction networks, the most central proteins have been found to be most essential to biological processes, removal of which being more deleterious than that of less central proteins [Jeo+01]. Similar results have been found in gene interaction networks [Özg+08]. *Pathway analysis* in gene and protein interaction networks have been shown to find biological pathways, series of interactions among molecules in a cell that leads to a certain product or a change in a cell [SO+07]. *Cluster analysis* on protein interaction networks allows for the identification of *functional modules* — protein complexes which perform a singular function [CY06].

2.6 Concept and topic extraction

In the disciplines of text mining, information extraction, and topic modelling, extraction of keywords and topics from a corpus are common tasks. Over the years, these disciplines have developed a number of methods for these tasks, varying in precision and accuracy. We will shortly discuss some commonly used methods.

Keyword extraction is the task of extracting keywords from a body of text. A commonly used method for this is *tf-idf* (term frequency inverse document frequency) [SJ72], which compares the frequency of a term within a single document to the frequency of that term within the entire corpus. Terms which occur frequently in the document but infrequently in the entire corpus obtain a high score, and are deemed more relevant to that document — these are keywords.

Topic extraction is the task of clustering terms into topics. Two commonly used methods for this task are *LSI* (latent semantic indexing) [Dee+90] and *LDA* (latent Dirichlet allocation) [Ble+03]. Both methods leverage term co-occurrence within documents to determine the likelihood of them being in the same topic, yielding a set of topics each consisting of a set of terms. A downside to these methods is that they do not scale well to corpora with many terms or documents, although a distributed adaptation to LDA has been proposed to mitigate this issue [Wan+09].

3 Concepts

> *Darwin's 'survival of the fittest' is really a special case of a more general law of survival of the stable. The universe is populated by stable things. A stable thing is [...] permanent enough or common enough to deserve a name.*
>
> — **Richard Dawkins**
> Evolutionary Biologist, Author of 'The Selfish Gene'

3.1 Introduction

In this Chapter, we will introduce some key concepts and definitions that will be used throughout the remainder of the thesis. Some of these concepts are part of common memetic theory, while others are novel contributions derived from existing work. First, we discuss our own theoretical contribution to the process of meme identification through tracing of their replication process in Section 3.2; we will then move on to a discussion of meme interactions, and show how a similar method to meme identification can also aid in the identification of meme interactions, in Section 3.3; and finally, we will conclude this Chapter by summarizing our theoretical contributions, and how they can be evaluated, in Section 3.4. Derivations of the symbols introduced in this Section can be found in Appendix A.

3.2 Identification of memes

In this Section, we will build upon the related work discussed in Chapter 2 to arrive at a general method for the identification of memes, which we will call the Memeseeker method. This method leverages one of the innate characteristics of memes as replicators — their replication process through imitation — to identify these memes in data on human correspondence, i.e. the things people read and write.

The Memeseeker method is inspired by the works of Kuhn et al. [Kuh+14] and Christakis and Fowler [CF07; CF08; CF13], which were discussed in Section 2.3. The method roughly consists of three steps, which we distilled from the work by Kuhn et al.: first, we construct a network over which memes could replicate — a *propagation network* — discussed in Section 3.2.2; then, we quantify how much of the occurrence of a particular behavioral pattern could be explained by such memetic replication, by assigning each behavioral pattern a *propagation score*, described in Section 3.2.3; and finally, we select those behavioral patterns for which the propagation score significantly deviates from what is expected were that pattern not a meme, a process we call *meme selection*, discussed in Section 3.2.4. First, however, we will start off by discussing the replication process of memes in more detail.

3.2.1 Meme replication

A meme is whatever it is that is passed on by imitation [Bla00]. As such, in order to find memes, we must find that which is imitated. Here, imitation should be regarded in the broad sense of the word, and may happen even though the imitator and imitatee are completely unaware of the process. Imitation happens when one individual observes the behavior of another individual, and proceeds to engage in the same behavior. In memetic parlance, a meme replicates when an *ego* is *exposed* to that meme, encoded in an *artefact*, and subsequently *expresses* that meme themselves by encoding it into another artefact. We will further explain these terms below. The process of meme replication through imitation is illustrated in Figure 3.1.

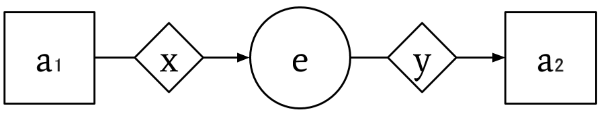

Fig. 3.1.: Schematic of meme replication through imitation. Here, ego e is exposed to a meme encoded in artefact a_1 through exposure x, and subsequently expresses it into artefact a_2 through expression y, completing the cycle.

Artefacts and egos In memetic theory [Bla00], an *artefact* is any intermediate form a meme can take, as long as some information about the original behavior is retained, and can thus be decoded. In this thesis, however, we restrict ourselves to *textual artefacts* — i.e. pieces of text — in which memes take the form of ngrams of length 1 through 4. The behavioral pattern in this case is the act of writing about the concept denoted by such an ngram.

An *ego* is any agent capable of interpreting memes encoded in artefacts, and engaging in the same behavior by re-encoding them into other artefacts. In this thesis, an ego is any agent able to decode the behavior denoted by an ngram — the behavior of writing the concept that the ngram denotes — and able to engage in the same behavior — writing about the concept themselves. Memetic theory considers both artefacts and egos as *vehicles* which *carry* memes [Bla00].

Let $A = \{a_1, a_2, ..., a_{|A|}\}$ be the set of all artefacts, and let some artefact $a \in A = \{m_1, m_2, ..., m_{|a|}\}$ represent the set of memes $m_1, m_2, ..., m_{|a|}$ that artefact a carries. Furthermore, let $E = \{e_1, e_2, ..., e_{|E|}\}$ be the set of all egos, and let some ego $e \in E = \{a_1, a_2, ..., a_{|e|}\}$ represent the set of artefacts $a_1, a_2, ..., a_{|e|} \in A$ that ego e expressed. An ego is said to carry all memes encoded in the artefacts it expressed.

Exposure and expression Exposure is the act of some ego encountering the memes carried by some artefact at some point in time. Expression is the act of some ego creating some artefact at some point in time, through encoding some memes into that artefact.

Let $X = \{x_1, x_2, ..., x_{|X|}\}$ with $x_1, x_2, ..., x_{|X|} \in A \times E \times \mathbb{R}$ be the set of all exposures, each of which of some ego $e \in E$ to some artefact $a \in A$ at some real-valued time $t \in \mathbb{R}$. Similarly, let $Y = \{y_1, y_2, ..., y_{|Y|}\}$ with $y_1, y_2, ..., y_{|Y|} \in A \times E \times \mathbb{R}$ be the set of all expressions, each of which by some ego $e \in E$ of some artefact $a \in A$ at some real-valued time $t \in \mathbb{R}$.

Furthermore, given some exposure or expression $z \in X \cup Y$, let $z^{(a)} \in A$ denote the artefact involved in z, let $z^{(e)} \in E$ denote the ego involved in z, and let $z^{(t)} \in \mathbb{R}$ denote the time at which z occurred.

3.2.2 Propagation network

In order to trace down those behavioral patterns that are due to imitation, we must establish potential imitation pathways over which behavioral patterns could have propagated. We call the network of potential imitation pathways a *propagation network*. Narrowing down such propagation pathways allows us to trace propagation on a web scale, as we can exclude pathways over which memes could not have propagated. In the citation networks used by Kuhn et al. [Kuh+14], authors of publications explicitly denote upon which publications they base their own work, and thus the citation network can be regarded as a propagation network. In the works by Christakis and Fowler [CF07; CF08; CF13], the propagation network is formed through inferring social ties between the subjects of study, and thus being able to trace which egos were influenced by which other egos at specific moments in time. Let us now introduce the concept of *influence*, after which we will show how to use this concept to construct a propagation network.

Influence In order to find that which is imitated, we introduce the notion of influence. We say that there is an influence from some artefact on another when some meme could have propagated from the former to the latter artefact. Given two artefacts $a_1, a_2 \in A$, there is an influence of a_1 on a_2 — denoted as $a_1 \rightarrow a_2$ — if and only if there was some ego $e \in E$ that was exposed to a_1 before expressing a_2.

$$a_1 \rightarrow a_2 \Leftrightarrow (\exists x \in X; \exists y \in Y)[x^{(a)} = a_1 \wedge y^{(a)} = a_2 \wedge x^{(e)} = y^{(e)} \wedge x^{(t)} < y^{(t)}]$$

Propagation network We can now construct a network of potential imitation pathways over which memes can propagate, which we will call a propagation network, defined as $G_{propagation} = (A, \{(a_1, a_2) \in A \times A : a_1 \rightarrow a_2\})$. Here, each artefact in A is represented by a node, with an edge from $a_1 \in A$ to $a_2 \in A$ when $a_1 \rightarrow a_2$.

3.2.3 Propagation score

After construction of the propagation network, we now can analyze how much the occurrence of a particular behavioral pattern can be explained by memetic propagation — its "meme-ness". We quantify this "meme-ness" through the construction of a *propagation score*, which measures the effect that prior influence of a pattern has on the subsequent occurrence of that pattern. For each pattern, we will analyze whether or not a vehicle (artefact or ego) is a carrier of that pattern, and whether or not it was influenced by that pattern. As we will discuss next, we can differentiate these vehicles into four groups: *sticking*, *sparking*, *non-sticking*, and *non-sparking* vehicles, which allows us to determine the effect of prior influence through the analysis of so-called *contingency tables*.

Sticking and sparking Sticking and sparking were first (informally) defined by Kuhn et al. [Kuh+14], though only for artefacts. Here, we establish formal definitions for these terms. For some meme m, we consider a vehicle to be *sticking* m if it carries m and we detect some a prior influence of m — imitation. If a vehicle carries m but we do *not* detect some prior influence of m, the vehicle is *sparking* m — invention. Furthermore, we distinguish between vehicles not carrying m for which there *is* a prior influence of m — *non-sticking* — and for which there is *not* a prior influence of m — *non-sparking*.

Propagation vehicle We will now define sticking and sparking for both types of vehicles. An artefact $a \in A$ is sticking meme m ($a \in A_m^m$) if it was influenced by an artefact carrying m, and carries m itself; a is sparking m ($a \in A_m^{\cancel{m}}$) if it was not influenced by an artefact carrying m, but does carry m itself; a is non-sticking m ($a \in A_{\cancel{m}}^m$) if it was influenced by an artefact carrying m, but does not carry m itself;

and finally, a is non-sparking m ($a \in A_{\not{m}}^{\not{m}}$) if it was neither influenced by an artefact carrying m nor carries m itself.

Similarly, an ego $e \in E$ is sticking m ($e \in E_m^m$) when it was exposed to m before the first expression of m; e is sparking m ($e \in E_m^{\not{m}}$) when it was not exposed to m before the first expression of m; e is non-sticking m ($e \in E_{\not{m}}^m$) when it was exposed to m but did not express m; and finally e is non-sparking m ($e \in E_{\not{m}}^{\not{m}}$) when it was neither exposed to nor expressed m.

Contingency tables Measuring the effect that some condition has on some outcome is a common task within clinical and epidemiological studies. Often, this results in the analysis of so-called *contingency tables*, which are displayed for artefact and ego populations in Table 3.1 and Table 3.2, respectively. In both tables, the amount of sticking, non-sticking, sparking, and non-sparking vehicles are assigned to the same respective variables, which we can then use to calculate the effect for either vehicle. In Chapter 4, we will evaluate which vehicle performs best with regards to meme identification.

	carrying m	not carrying m	Total
influenced by m	$p_m = \lvert A_m^m \rvert$	$q_m = \lvert A_{\not{m}}^m \rvert$	$p_m + q_m$
not influenced by m	$r_m = \lvert A_m^{\not{m}} \rvert$	$s_m = \lvert A_{\not{m}}^{\not{m}} \rvert$	$r_m + s_m$
Total	$p_m + r_m$	$q_m + s_m$	$p_m + q_m + r_m + s_m$

Tab. 3.1.: Contingency table for behavioral pattern m within the artefact population. Here, p_m denotes the amount of artefacts sticking behavioral pattern m; q_m the amount of artefacts non-sticking m; r_m the amount of artefacts sparking m; and s_m the amount of artefacts non-sparking m.

	carrying m	not carrying m	Total
influenced by m	$p_m = \lvert E_m^m \rvert$	$q_m = \lvert E_{\not{m}}^m \rvert$	$p_m + q_m$
not influenced by m	$r_m = \lvert E_m^{\not{m}} \rvert$	$s_m = \lvert E_{\not{m}}^{\not{m}} \rvert$	$r_m + s_m$
Total	$p_m + r_m$	$q_m + s_m$	$p_m + q_m + r_m + s_m$

Tab. 3.2.: Contingency table for behavioral pattern m within the ego population. Here, p_m denotes the amount of egos sticking behavioral pattern m; q_m the amount of egos non-sticking m; r_m the amount of egos sparking m; and s_m the amount of egos non-sparking m.

Propagation effect Given the contingency tables as above, we find there are three commonly used measures to quantify the effect of exposure on expression: the *absolute risk* (AR) of exposure on expression; the *relative risk* (RR) of exposure on expression, used by Kuhn et al. [Kuh+14]; and the *odds ratio* (OR), used by Christakis and Fowler. See the equations below for their definitions, taken from Agresti [Agr07]. Through our evaluations in Chapter 4, we will investigate which of these effect metrics performs best, and should be used as a propagation score.

$$AR(m) = \frac{p_m}{p_m + q_m} - \frac{r_m}{r_m + s_m}$$

$$RR(m) = \frac{p_m}{p_m + q_m} / \frac{r_m}{r_m + s_m}$$

$$OR(m) = \frac{p_m}{r_m} / \frac{q_m}{s_m}$$

3.2.4 Meme selection

Patterns with a low number of occurrences have a relatively high probability of obtaining a high propagation score, due to a large amount of variance. To counter such noise, Kuhn et al. multiply each pattern's propagation score by the relative frequency of that pattern in the corpus. However, this has the unintended consequence that some patterns which occur very often — i.e. stopwords — do obtain a high combined score. Indeed, Kuhn et al. filter the most frequent patterns out of their results to get rid of these stopwords. Here, we propose an alternative, more general method for the selection of memes.

Null hypotheses of no effect Inspired by analyses of microarray data in genomics, we only select behavioral patterns for which the effect is significant. For each pattern m we set up a null hypothesis H_0^m of no effect, with the alternative hypothesis H_a^m that there is an effect, as below:

$$\begin{aligned} H_0^m : \frac{p_m}{p_m + q_m} &= \frac{r_m}{r_m + s_m} \\ \Leftrightarrow AR(m) &= 0 \\ \Leftrightarrow RR(m) &= 1 \\ \Leftrightarrow OR(m) &= 1 \end{aligned}$$

$$\begin{aligned} H_a^m : \frac{p_m}{p_m + q_m} &\neq \frac{r_m}{r_m + s_m} \\ \Leftrightarrow AR(m) &\neq 0 \\ \Leftrightarrow RR(m) &\neq 1 \\ \Leftrightarrow OR(m) &\neq 1 \end{aligned}$$

In other words, the null hypothesis for a given pattern is that it does not behave like a meme, while the alternative hypothesis is that it does behave like a meme. We can then select those behavioral patterns for which we can reject the associated null hypothesis. Since the null hypothesis is the same for each choice of effect, only the choice of vehicle — egos or artefacts — will influence selection.

Hypothesis testing We test our null hypotheses using the Wald test for difference in binomial proportions [Agr07]. We use the adjustment for this interval as proposed in [AC00], which should be more powerful than the unadjusted version when some cell counts in the contingency table are low, which we expect to happen due to many ngrams extracted from natural language only occurring infrequently [MS99].

As is common with multiple hypothesis testing [BH95], we correct for the false discovery rate by modifying the acquired P-values with the Benjamini-Yekutieli procedure [BY01]. We use the Benjamini-Yekutieli procedure instead of the more commonly used Benjamini-Hochberg procedure [BH95], as the former is more robust against dependent hypotheses, and memetic theory predicts memes to be dependent upon one another. We then reject all null hypotheses with a modified P-value < 0.05, and select those patterns with non-rejected null hypotheses, such that we are left with all patterns where there is a significant propagation score — memes.

3.3 Identification of meme interaction

In this Section, we will discuss how we could potentially find the interactions that memetics predicts, through a slight alteration of the Memeseeker method we introduced in the previous Section. This alteration allows us to quantify how much exposure to some pattern affects expression of some *other* pattern. We will call this metric the *co-propagation score*, as it quantifies the tendency of two memes to propagate together through the propagation network. We will further discuss this co-propagation score for the quantification of meme interactions, which can be leveraged for the construction of *meme interaction networks*, in Section 3.3.1; and then, we will discuss how the use of centrality metrics applied to meme interaction networks should help us identify key memes in Section 3.3.2.

3.3.1 Co-propagation score

A measure of meme co-propagation, defined as the effect that exposure to some meme i has on the probability of expression of some other meme j, is taken as a method for the identification of meme interaction. The method we propose has one main benefit over the method proposed by Best [Bes97], discussed in Section 2.4, who derives the cross-correlation between memes occurrence frequencies: by leveraging the propagation network, our method is independent of time spacing between meme occurrence, only looking at time order, which allows us to more accurately identify meme interactions. Furthermore, we see the conceptual similarity between the methods for meme identification and interaction identification as an added benefit.

Through the evaluations of the Memeseeker method in Chapter 4, we will find that the usage of egos as vehicle, combined with absolute risk as propagation effect, yields the best results. Therefore, we will no longer consider artefact vehicles, or relative risk and odds ratio effects measures.

Let us expand definitions of sticking, non-sticking, sparking, and non-sparking vehicle populations, as we defined them in Section 3.2.3, to take into account differing memes for influence and carrying. An ego $e \in E$ is sticking j through i ($e \in E_j^i$) when it was exposed to i before the first expression of j; e is sparking j through i ($e \in E_j^{\not{i}}$) when it was not exposed to i before the first expression of j; e is non-sticking j through i ($e \in E_{\not{j}}^i$) when it was exposed to i but did not express j; and finally e is non-sparking j through i ($e \in E_{\not{j}}^{\not{i}}$) when it was neither exposed to i to nor expressed j. We display the ego populations in a contingency table, shown in Table 3.3.

	carrying j	not carrying j	Total
influenced by i	$p_{ij} = \lvert E_j^i \rvert$	$q_{ij} = \lvert E_{\not{j}}^i \rvert$	$p_{ij} + q_{ij}$
not influenced by i	$r_{ij} = \lvert E_j^{\not{i}} \rvert$	$s_{ij} = \lvert E_{\not{j}}^{\not{i}} \rvert$	$r_{ij} + s_{ij}$
Total	$p_{ij} + r_{ij}$	$q_{ij} + s_{ij}$	$p_{ij} + q_{ij} + r_{ij} + s_{ij}$

Tab. 3.3.: Contingency table for the interaction by meme i on meme j.

We can then use these contingency tables to measure the strength of each interaction by calculating the *absolute risk* (AR) that influence to meme i has on the probability of carrying meme j: $AR(i,j) = \frac{p_{ij}}{p_{ij}+q_{ij}} - \frac{r_{ij}}{r_{ij}+s_{ij}}$ [Agr07], allowing us to quantify the influence that each meme has on each other meme. We calculate these the co-propagation scores for every meme on every other meme. As the number of calculations scales quadratically with the number of memes, we only perform it for patterns which were identified as memes through the Memeseeker method, instead of performing it for every pattern. We then filter out insignificant interactions using the adjusted Wald test for binomial proportions, in the same fashion as was discussed in Section 3.2.4.

We construct a meme interaction network by linking interacting memes with directed, weighted edges denoting the strength and direction of the interaction. A potential use for such networks will be discussed next.

3.3.2 Meme centrality

The Memeseeker method quantifies how much some textual pattern propagates like a meme within a community — its propagation score — and in the evaluation in Chapter 4, we use this quantification to rank memes. Through those evaluations, we find support for the intuition that the higher the propagation score of a meme, and thus the more it propagates like a meme, the higher its perceived relevance is to the corresponding field of study.

Here, we argue that not only a meme's propagation score, but also how it relates to other memes determines its perceived relevance. In particular, we hypothesize that memes which interact more strongly with other memes will be perceived to be more relevant. Therefore, some measure of a meme's centrality within an interaction network should be a good quantification for its perceived relevance, i.e. the more a meme interacts with other memes, the more central it will be in the interaction network, and the more relevant it will be perceived to be. We will evaluate this hypothesis in Chapter 5.

3.4 Conclusion

We have introduced the Memeseeker method, which, through construction of a *propagation network*, should allow us to quantify the "meme-ness" of behavioral patterns by assigning each pattern a *propagation score*. This score captures the effect that *influence* of a pattern has on a vehicle *carrying* of that pattern. We then *select* those patterns for which this effect is significant, and regard them as memes. We identified two types of *vehicles* — *egos* and *artefacts* — for which we can quantify an effect using three methods commonly used in clinical studies — *absolute risk*, *relative risk*, and the *odds ratio*. We will evaluate this method, and which of the aforementioned choices for vehicle and effect allow us to best identify memes, in Chapter 4.

We further showed how we should be able to identify the *interactions* between memes, through quantification of the effect that prior influence of some meme has on the probability of a vehicle carrying some *other* meme. This then allows us to construct a *meme interaction network*, through which emergent properties — e.g. meme *centrality* — of interacting memes can be derived. We will evaluate this method for the identification of meme interactions, through evaluating whether meme centrality in the constructed meme interaction network allows us to better identify key concepts, in Chapter 5.

4 Evaluation of meme identification

> *Words are memes that can be pronounced.*
>
> — **Daniel C. Dennett**
> Philosopher, Author

4.1 Introduction

The Memeseeker method identifies memes through a three-step process, as discussed in Section 3.2: first, we construct a network of potential meme propagation pathways; then, we quantify how much each behavioral pattern propagates over this network; and finally, we select those patterns as memes for which their propagation is significantly like a meme.

In this Chapter, we will evaluate the Memeseeker method for the identification of memes. In these evaluations we will restrict ourselves to *scientific memes* only, i.e. memes represented as text — particularly, ngrams of length 1 through 4 — denoting important topics and concepts to particular scientific fields. Note that such scientific memes are a valid subset of memes: the behavior that is being imitated as the meme propagates from person to person is the behavior of writing about a particular scientific topic or concept.

To evaluate our method of meme identification, and thus to be able to address the first research question ("Can we identify memes propagating through online communities?"), we apply the method to five science-oriented communities on the online platform Reddit, each organized around a specific field of study — aerospace engineering, computer science, data science, medicine, and psychology. We choose to evaluate our method on science-oriented correspondence as we have an intuition for what kind of memes we should expect to detect in this correspondence, namely important concepts and topics within the field. Besides, this gives us some comparison to the method proposed by Kuhn et al. [Kuh+14], which was also applied to science-oriented correspondence.

Using correspondence data extracted from these five science-oriented communities, we will first perform a coarse, large-scale offline evaluation in Section 4.2, followed by a finer, but smaller-scale study with human participants in 4.3. We will discuss the results of these evaluations with respect to the three steps of the Memeseeker method in Section 4.4. In Chapter 6 we will discuss the results with respect to the research questions. A replication of the offline evaluation on data from the question and answer platform StackExchange can be found in Appendix C.

4.2 Offline evaluation

The goal of the offline evaluation is to get a high-level insight into the general performance of the Memeseeker method. This will ensure that our method performs at least well enough to conduct a more fine-grained analysis later, and to allow us to understand the general performance of the method by analyzing the entirety of the selected memes.

4.2.1 Procedure

We evaluate our method in a similar fashion as information retrieval problems, where we rank the selected memes by their propagation score, and use relevancy-based metrics to quantify the performance of the system. In this offline evaluation, we will use an automated procedure to assign relevancy to each selected meme. We say that an ngram is *relevant* when it is the title of a Wikipedia article, and not a stopword — as judged by the Python Natural Language Toolkit [Bir+09]. We realize this metric is broad, however it can be performed at scale. A more fine-grained analysis will be conducted later.

Dataset We acquire five datasets, each containing all comments created within a specific science-oriented community: aerospace (containing discussions around aerospace engineering); compsci (computer science); datascience (data science); medicine (medicine); psychology (psychology). We obtain these datasets from a publicly available dump of Reddit comments on Google BigQuery. We picked a diversity of datasets to evaluate whether the process works equally across domains. See Table 4.1 for a summary of the datasets.

Text cleaning We perform a simple cleaning process on the textual contents of the comments. We remove double quotes ("), periods (.), commas (,), exclamation marks (!) and question marks (?). All text is converted into lowercase. Stopwords are not removed.

Egos, artefacts and behavioral patterns Per dataset, each comment represents a single artefact, which was expressed by an ego — the user who created it. The behavioral patterns from which we will be selecting and scoring memes are all unique ngrams of length 1 through 4, extracted from the textual contents of the comments. A summary of the data, displaying the number of egos, artefacts, and unique ngrams per dataset can be found in Table 4.1.

community	# egos	# artefacts	# ngrams
aerospace	6,949	27,130	2,653,815
compsci	55,563	216,554	17,284,739
datascience	13,050	59,265	5,666,345
medicine	84,719	493,871	34,701,719
psychology	115,781	348,659	24,630,907

Tab. 4.1.: Summary of datasets, showing the number of egos, artefacts and unique ngrams for each dataset.

Constructing the propagation network In order to construct a propagation network for our Reddit datasets, we must determine which ngrams an ego has expressed, and to which ngrams an ego was exposed. Establishing expression is trivial, as each artefact expresses the ngrams it contains. However, as the datasets do not contain any information on which users have read which comments, establishing exposure is less trivial, and requires us to develop a heuristic.

We represent the correspondence by users on Reddit as a forest of tree-structured comment threads, as displayed in Figure 4.1. A comment's ancestors are those comments which are higher up this hierarchy. We heuristically assume that a user leaving a comment was exposed to all ancestors of that comment. As such, we can construct a propagation network by connecting each user's comments to their ancestors, and to the ancestors of comments that were created earlier by that user.

4.2.2 Independent variables

Propagation vehicle In Section 3.2.3, we discussed how memetic propagation can be detected as memes propagate through *vehicles*: either *egos* or *artefacts*. We will investigate how the choice of propagation vehicle affects ranking and precision.

Propagation effect In Section 3.2.3, we discussed how memetic propagation can be quantified through three types of effect metrics: *absolute risk*, *relative risk* and the *odds ratio*. We will investigate how the choice of propagation effect affects ranking.

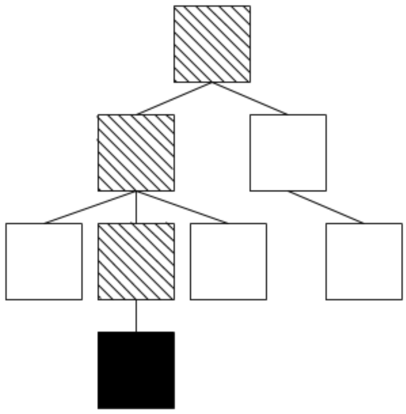

Fig. 4.1.: Example of the hierarchy of correspondence on Reddit. Discussion on Reddit happens within threads, where users can comment on each other's comments. Here, we visualize the discussion within such a thread, where each square represents a comment by a different user. We heuristically assume that a user is exposed to all ancestors of the comments they create. In this image, the user creating the comment marked in black is assumed to have been exposed to the shaded comments.

4.2.3 Dependent variables

We will evaluate the metrics described below at rank 100 — to compare to the online evaluation — and at the final rank of each generated list (N).

(Mean) normalized discounted cumulative gain To investigate how the choice of vehicle and effect affect *ranking*, we will measure the normalized discounted cumulative gain (nDCG), a commonly used metric for evaluating the performance of ranking. We will summarize the nDCG by taking the mean across datasets, such that we can compare the choice of vehicle and effect.

(Mean) precision To investigate how choice of vehicle affects *selection*, we will measure the precision, the fraction of memes that are deemed relevant. We will summarize the precision by taking the mean across datasets, such that we can compare the choice of vehicle.

4.2.4 Hypotheses

Ranking and selection perform best for egos, compared to artefacts As an ego can produce many artefacts containing a given meme — simply reproducing one artefact many times will be sufficient, and can be performed quickly and cheaply in the digital age — measurement of artefact vehicles may not reliably reflect the true "meme-ness" of an ngram. We should therefore expect measurement of egos to yield better results, and to outperform artefacts both in terms of both ranking and selection, which should be reflected in the nDCG and precision scores, respectively.

Ranking performs best for absolute risk, compared to relative risk and odds ratio As stated in Schechtman [Sch02], the absolute risk has the advantage of both reflecting the prior probability of expressing a certain ngram and the increase in probability after exposure to that ngram, which may be more advantageous for more commonly occurring ngrams. On the other hand, relative risk will tend to give a lot of weight for even a minor increase in probability, given that the prior probability was already low, and thus yield a higher score for less commonly occurring ngrams. For small prior probabilities, the odds ratio will yield similar values to the relative risk, but will tend to exaggerate the relative risk for higher prior probabilities [Vie08]. All in all, we should expect the absolute risk to outperform both relative risk and odds ratio in terms of ranking, which should be reflected in the nDCG scores.

4.2.5 Results

We compare the results for vehicle (ego or artefact) and effect (absolute risk, relative risk, or odds ratio); resulting in 6 conditions. Comparisons are given for five datasets (aerospace, compsci, datascience, medicine, psychology). The results of the evaluation are summarized per dataset in Table 4.2, and a summary across datasets is given in Table 4.3. A small manual selection of memes, as ranked and selected through ego vehicles and absolute risk, can be found in Table 4.4. For the top 100 memes, as ranked and selected through ego vehicles and absolute risk, we refer the reader to Appendix B.

Ranking and selection perform best for egos, compared to artefacts We observe that measurement of egos significantly outperforms measurement of artefacts, both in terms of nDCG and precision. In Figure 4.2 we show the progression of precision over rank, which shows low precision of artefact scores at low ranks, only to increase drastically for higher ranks. This result is also visible in Table 4.2 and Table 4.3, where we observe that precision@100 values are extremely low for artefacts. Manual inspection of the generated rankings shows this is largely caused by a high prevalence of repeated comments, mostly created by bots, but often times also created by moderators of the communities — for example, comments reminding users of the community's rules.

Ranking performs best for absolute risk, compared to relative risk and odds ratio As expected, absolute risk does outperform relative risk and odds ratio in terms of nDCG, and relative risk and odds ratio do appear to perform similar. However,

community	vehicle	effect	# memes (N)	nDCG @100	nDCG @N	precision @100	precision @N
aerospace	egos	AR	520	0.977	0.982	0.970	0.681
aerospace	egos	RR	520	0.960	0.983	0.950	0.681
aerospace	egos	OR	520	0.961	0.983	0.950	0.681
aerospace	artefacts	AR	1,429	0.009	0.742	0.010	0.421
aerospace	artefacts	RR	1,429	0.008	0.749	0.010	0.421
aerospace	artefacts	OR	1,429	0.000	0.748	0.000	0.421
compsci	egos	AR	5,923	0.761	0.954	0.840	0.514
compsci	egos	RR	5,923	0.692	0.936	0.750	0.514
compsci	egos	OR	5,923	0.660	0.935	0.750	0.514
compsci	artefacts	AR	9,218	0.000	0.838	0.000	0.421
compsci	artefacts	RR	9,218	0.000	0.835	0.000	0.421
compsci	artefacts	OR	9,218	0.000	0.835	0.000	0.421
datascience	egos	AR	1,399	0.937	0.968	0.930	0.631
datascience	egos	RR	1,399	0.905	0.959	0.910	0.631
datascience	egos	OR	1,399	0.906	0.960	0.910	0.631
datascience	artefacts	AR	2,981	0.000	0.792	0.000	0.438
datascience	artefacts	RR	2,981	0.000	0.794	0.000	0.438
datascience	artefacts	OR	2,981	0.009	0.794	0.010	0.438
medicine	egos	AR	16,316	0.949	0.958	0.950	0.410
medicine	egos	RR	16,316	0.819	0.933	0.810	0.410
medicine	egos	OR	16,316	0.831	0.935	0.820	0.410
medicine	artefacts	AR	15,269	0.000	0.862	0.000	0.435
medicine	artefacts	RR	15,269	0.000	0.857	0.000	0.435
medicine	artefacts	OR	15,269	0.000	0.857	0.000	0.435
psychology	egos	AR	8,262	0.924	0.961	0.910	0.475
psychology	egos	RR	8,262	0.885	0.939	0.860	0.475
psychology	egos	OR	8,262	0.885	0.940	0.860	0.475
psychology	artefacts	AR	19,985	0.000	0.848	0.000	0.328
psychology	artefacts	RR	19,985	0.000	0.833	0.000	0.328
psychology	artefacts	OR	19,985	0.000	0.833	0.000	0.328

Tab. 4.2.: Results of the offline evaluation. This table shows the nDCG and precision, evaluated at rank 100 and at the final rank N (the total number of selected memes) for each combination of vehicle and effect, per dataset.

vehicle	effect	mean nDCG @100	mean nDCG @N	mean precision @100	mean precision @N
egos	AR	0.910	0.965	0.920	0.542
egos	RR	0.852	0.950	0.856	0.542
egos	OR	0.848	0.951	0.858	0.542
artefacts	AR	0.002	0.816	0.002	0.409
artefacts	RR	0.002	0.814	0.002	0.409
artefacts	OR	0.002	0.813	0.002	0.409

Tab. 4.3.: Summary of the offline evaluation. This table shows the mean nDCG and mean precision across datasets, evaluated at rank 100 and at the final rank N, for each combination of vehicle and effect.

seeing as the difference between the performance of absolute risk, relative risk

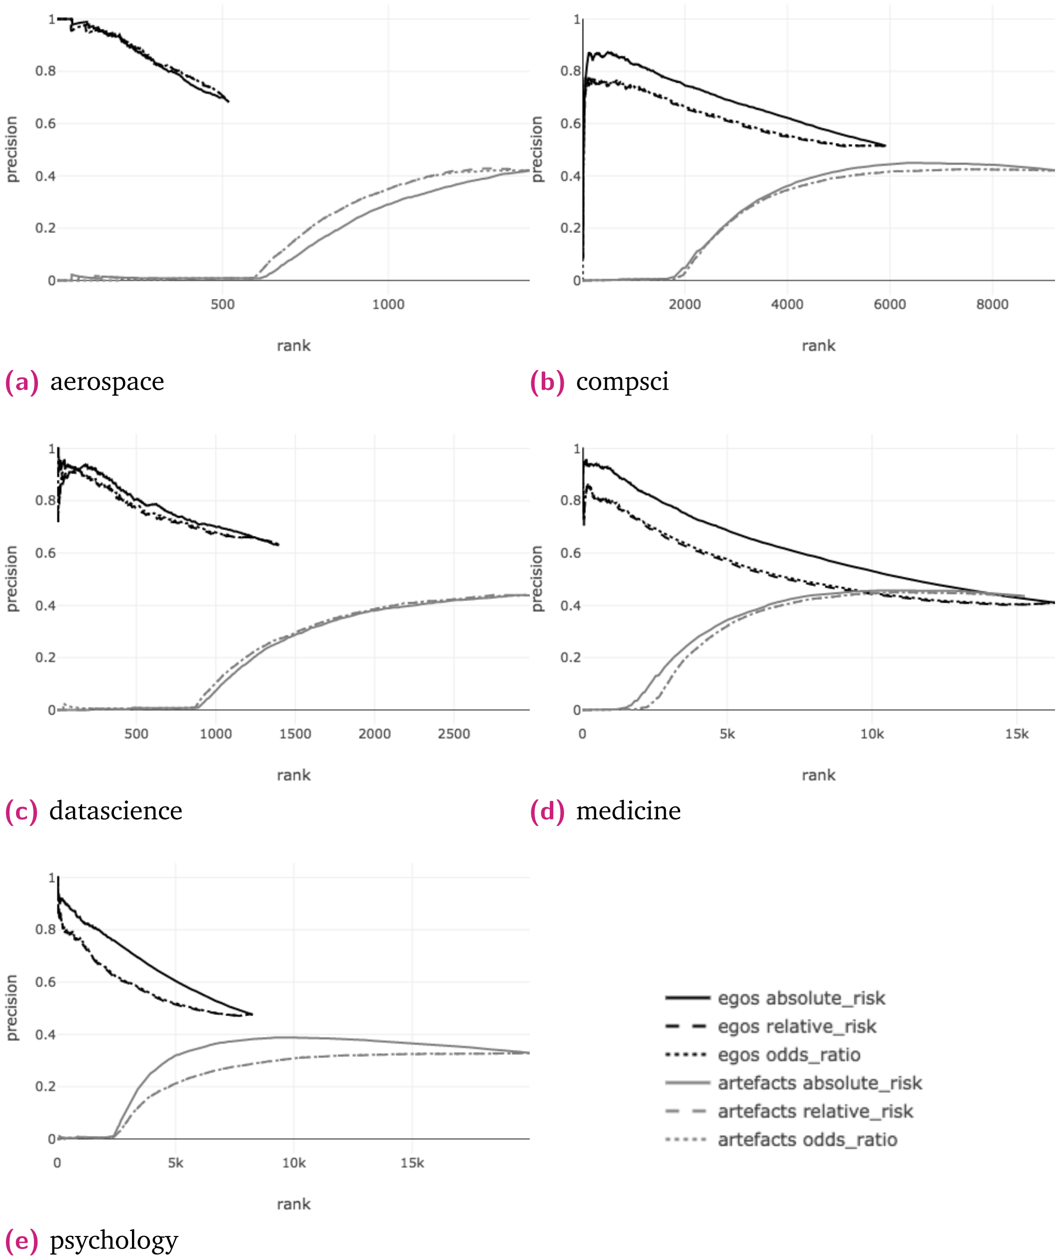

(a) aerospace (b) compsci (c) datascience (d) medicine (e) psychology

Fig. 4.2.: Offline evaluation: rank vs. precision for each combination of vehicle and effect, per dataset. Here, black lines denote ego vehicles, grey lines denote artefact vehicles. Solid lines denote absolute risk, dashed lines denote relative risk, and dotted lines denote odds ratio. As relative risk and odds ratio perform very similar, they overlap in these figures.

and odds ratio is relatively small, we will further investigate it through our online evaluation.

community	rank	propagation score	meme
aerospace	13	0.267	airfoil
aerospace	29	0.228	thrust
aerospace	35	0.220	propulsion
compsci	18	0.391	pumping lemma
compsci	21	0.307	cache misses
compsci	34	0.272	hash
datascience	1	0.472	python 2
datascience	13	0.250	clusters
datascience	35	0.219	random forest
medicine	11	0.292	dengue
medicine	24	0.259	ebola
medicine	32	0.252	bicarb
psychology	9	0.313	emdr
psychology	27	0.249	mbti
psychology	41	0.224	p-value

Tab. 4.4.: A manual selection of memes for each dataset, for illustrative purposes. These memes were selected from the top 100 memes as ranked and selected through ego measurement and absolute risk. The entire list can be found in Appendix B.

4.3 Online evaluation

The offline evaluation gives a good sense of the relevance of the topics already. However, this is limited by the concepts that currently are listed in Wikipedia. In order to evaluate the degree of relevance in a more granular fashion, we also asked experts in each scientific field to judge the relevance of the found memes with respect to their field. This research has been reviewed by TU Delft's Human Research Ethics Committee.

4.3.1 Procedure

For each dataset and associated field, we asked a number of participants to judge the relevancy of selected memes. The ranking is calculated in a similar way as for the offline evaluation. However, due to the poor performance of artefact measurement, we will drop it from our online evaluation, only studying the ego vehicles.

Relevancy is judged by human expert annotators. To keep the evaluation manageable for participants, rankings will be truncated to a length of 100 for each effect type. Selected memes for each effect — absolute risk, relative risk, odds ratio — are combined into one list per field, filtering out duplicates. For each participant, the resulting list is randomized, and memes are presented one-by-one. For each meme, participants asked to denote its relevancy to their field of study on a 4-point scale: *0)* irrelevant; *1)* marginally relevant; *2)* relevant; and *3)* highly relevant. If participants did not recognize a meme, they were asked to look it up in order to judge the relevancy. If they could not find the term easily, they were asked to denote it as irrelevant.

4.3.2 Participants

Sixteen participants completed the study, across five disciplines: aerospace (2), computer science (6), data science (2), medicine (2); and psychology (4). Participants all completed at least a Bachelor-level formal education in their respective fields, with some participants holding a Masters or other post-graduate level degree.

4.3.3 Independent variables

Propagation vehicle We will only study *egos* in this online evaluation, since artefacts performed much worse than egos in the offline evaluation.

Propagation effect We will again investigate the influence of three kinds of propagation effect: absolute risk, relative risk, and odds ratio.

4.3.4 Dependent variables

(Mean) normalized discounted cumulative gain Since nDCG can be also applied to non-binary relevance metrics, we can use it to compare the ranking of each of our conditions in a more granular fashion. This allows us to better differentiate which of the effects performs better.

(Mean) precision Precision can only be applied to a binary relevance, while our evaluation task requires participants to judge relevance on a 4-level scale. In order to calculate precision, and thus to compare the results of the offline evaluation with the online evaluation, we need to convert the relevances to a binary scale. We say that an ngram is relevant when it is judged as at least marginally relevant by the participants, and irrelevant otherwise.

4.3.5 Hypotheses

We again expect the absolute risk to outperform relative risk and odds ratio. This is for the same reason as stated in the offline evaluation: the absolute risk yields a higher score for ngrams that occur more frequently, which should reduce the number of false-positives at the top of the list.

4.3.6 Results

The absolute risk again slightly outperforms the relative risk and odds ratio in terms of nDCG and precision (c.f., Tables 4.5 and 4.6).

community	effect	annotators (#)	mean nDCG	mean precision
aerospace	AR	2	0.844	0.840
aerospace	RR	2	0.824	0.845
aerospace	OR	2	0.824	0.845
compsci	AR	6	0.630	0.702
compsci	RR	6	0.634	0.610
compsci	OR	6	0.622	0.615
datascience	AR	2	0.790	0.835
datascience	RR	2	0.758	0.810
datascience	OR	2	0.775	0.815
medicine	AR	2	0.767	0.765
medicine	RR	2	0.781	0.705
medicine	OR	2	0.781	0.720
psychology	AR	4	0.547	0.525
psychology	RR	4	0.513	0.475
psychology	OR	4	0.513	0.472

Tab. 4.5.: Results of the online evaluation. This table shows the mean nDCG and mean precision, for each effect, per dataset. Here, we measure egos as vehicles.

effect	mean nDCG	mean precision
AR	0.716	0.733
RR	0.702	0.689
OR	0.703	0.694

Tab. 4.6.: Summary of the online evaluation. This table shows the mean of mean nDCG scores per dataset and the mean of mean precisions per dataset, for each effect. Here, we measure egos as vehicles.

4.4 Discussion

Here we discuss the implications of the results for each of the steps of the proposed method for meme extraction, as laid out in Section 3.2: construction of the propagation network, measurement of a propagation score quantifying a pattern's "meme-ness", and selection of those patterns that propagate significantly like memes.

4.4.1 Propagation network

Drawing upon the theoretical laid out in Section 3.2.2, we were able to construct a propagation network for data extracted from Reddit. For this dataset, comments took the form of *artefacts*, and their authors the form of *egos*. *Expression* occurs when a user creates a comment, and *exposure* occurs when we assume a user has read a comment. The behavioral patterns which constitute memes are the ngrams that occur in those comments. With the detection of exposure and expression we can infer imitation, and thus meme propagation, through the construction of a propagation network.

4.4.2 Propagation score

In Section 3.2.3, we identified two main components for the quantification of "meme-ness", i.e. the propagation score: the vehicle through which to track the propagation of memes, and the effect with which to quantify this propagation in terms of imitation.

Propagation vehicle We find that the choice of vehicle strongly influences the performance of propagation, with egos outperforming artefacts by a large margin, as is clearly visible in Table 4.3. In Figure 4.2 we observe that this is due to a large number of irrelevant memes obtaining a high rank, when measuring artefacts.

Even though we did expect egos to outperform artefacts, the *extent* by which they did is surprising. Through manual inspection of the results, we find that this difference is mostly due to non-human users — bots — "parroting" (near) identical messages many times. It appears to be unfeasible to exclude all comments created by bot users. Reddit bots are not formally registered, and while manually curated list of bots do exist, they are incomplete and out-of-date.

Through measurement of egos we can circumvent this issue, as this assigns equal influence to each ego, regardless of how many times that ego expresses a certain pattern. This ensures that repeated messages do not skew the results, and as such, provides a better insight into the actual propagation of memes. We therefore recommend the measurement of egos rather than artefacts, as long as the dataset contains enough egos to support a list of a specific rank and acceptable level of precision.

Propagation effect We test three methods commonly used in epidemiology to quantify the effect of some exposure on subsequent expression, and find that the choice of method impacts performance to a lesser extent. Through both Table 4.3 and Table 4.6 we observe that absolute risk systematically outperforms relative risk and the odds ratio, but only by a small margin. Through Figure 4.2 we do observe a noticeable difference between absolute risk on the one hand, and relative risk and odds ratio on the other, with the precision of absolute risk being higher than that of relative risk and odds ratio, for most ranks.

Besides performing better, we recommend absolute risk for reasons of interpretability: whereas the *relative risk* and odds ratio obtain values between 0 and infinity, with 1 denoting no effect, 0 denoting that expression only occurs without prior exposure, and infinity denoting that expression only occurs *with* prior exposure; *absolute risk* on the other hand is symmetric around 0, and obtains values which are bounded on both sides by -1 to 1. This symmetry allows us to better compare positive and negative impacts of expression on exposure.

4.4.3 Meme selection

Through Table 4.2, our offline evaluation finds that the selection process performs relatively well for egos, with a precision ranging from around 0.4 to almost 0.7 for all fields for the entire set of selected memes. Furthermore, our propagation scoring then ensures that the most relevant memes obtain the highest score, ranking them at the top of the resultant list, with the top 100 memes having a precision ranging from ca 0.7-0.9. This finding is further supported by our online evaluation, where we observe through Table 4.5 that the precision of our method ranges from ca 0.5-0.8 for the first 100 memes.

We further observe that the number of selected memes for ego measurement is higher than artefact measurement only for the medicine community, and that the precision eventually drops below that of artefact measurement (c.f., Figure 4.2d). This is unexpected given the results for the other four domains and communities. One possible explanation is a difference in the dataset compared to the other domains; with a high number of artefacts produced per ego (c.f., Table 4.1). Further study should elucidate this phenomenon.

5 Evaluation of meme interaction identification

> *The essence of any memeplex is that the memes inside it can replicate better as part of the group than they can on their own.*
>
> — **Susan Blackmore**
> Author of 'The Meme Machine'

5.1 Introduction

By quantifying the effect that influence by a meme has on the probability of a vehicle carrying the some *other* meme, we constructed a method for the identification of meme interactions in Section 3.3. As the identification of such meme interactions on such a fine scale is a novel task, direct evaluation of the correctness of our method is difficult. We will therefore perform an indirect evaluation. We assume that meme interaction networks are similar in function to other types of interaction networks, such as ecological networks and protein interaction networks. Under this assumption, certain analyses of the latter networks can be applied to meme interaction networks, yielding similar results. We discussed in Section 2.5 that centrality analysis applied to ecological networks and protein interaction networks has been shown to yield key species and proteins, respectively. Therefore, our evaluation involves testing the performance of ranking memes by their centrality in their meme interaction networks.

In this Chapter, we evaluate our meme interaction network against a baseline meme co-occurrence network, as was used in the work by Chavalarias and Cointet [CC13] discussed in Section 2.4. If ranking by centrality performs better in our meme interaction network than the meme co-occurrence network with respect to perceived relevance, we take this as support that our method identifies meme interactions better than simple co-occurrence. In Chapter 6 we will discuss the results with respect to the research questions.

5.2 Evaluation

In this Section, we will test the hypothesis that a meme's centrality within the meme interaction network is a better quantification of that meme's perceived relevance, compared to that meme's centrality in the meme co-occurrence network. Furthermore, we obtain centrality scores through a number of separate methods — betweenness centrality, eigenvector centrality, Katz centrality, PageRank, and HITS — such that we can find which method best quantifies the perceived relevance.

5.2.1 Procedure

We perform this evaluation for the top 100 memes found in Section 4.2, as selected and ranked by using egos as vehicles and absolute risk as effect measure, which can be found in Appendix B. We use the same datasets and human evaluation of perceived relevancy that were used for the evaluation of the Memeseeker method in Chapter 4.

Again, relevancy is judged by human expert annotators. For each meme, participants asked to denote its relevancy to their field of study on a 4-point scale: *0)* irrelevant; *1)* marginally relevant; *2)* relevant; and *3)* highly relevant. If participants did not recognize a meme, they were asked to look it up in order to judge the relevancy. If they could not find the term easily, they were asked to denote it as irrelevant.

5.2.2 Participants

As before, sixteen participants completed the study, across five disciplines: aerospace (2), computer science (6), data science (2), medicine (2); and psychology (4). Participants all completed at least a Bachelor-level formal education in their respective fields, with some participants holding a Masters or other post-graduate level degree.

5.2.3 Independent variables

In this evaluation, we have two independent variables: the type of network — meme interaction network versus meme co-occurrence network — and the centrality metric.

Network We will compare the meme interaction network to the meme co-occurrence network. For each dataset, we construct the meme interaction network as laid out in Section 3.3. The meme co-occurrence network is obtained through modelling each meme as a node, with edges denoting the co-occurrence of memes in artefacts, normalized using the Jaccard similarity measure [Sma73].

Centrality metric Throughout the years, many centrality metrics have been proposed, most of them focusing on particular domains or problems. We will calculate centrality scores using a few commonly used metrics which can be applied to weighted, directed networks. For each of these metrics, we use the Python implementation in NetworkX [Hag+08] using default parameters. The selected centrality metrics are discussed below.

Betweenness centrality [Fre77] is a measure of node centrality based on shortest paths in the network. The number of shortest paths between any two nodes that passes through a node determines the centrality of that node.

Eigenvector centrality [Bon87] is a measure of centrality based upon the calculation of eigenvectors within the network adjacency matrix, and is a measure of a node's influence in a network. In contrast to betweenness centrality, links to high-scoring nodes contribute more than links to low-scoring nodes.

Katz centrality [Kat53] was first proposed as a status index of individuals in social networks. It was developed as a generalization of eigenvector centrality, and by assigning each node a small initial centrality score, it will converge for networks for which eigenvector centrality will not find a solution (e.g. directed acyclic graphs).

PageRank [Pag+99] is based upon the idea of *random walks* through the network, using edge weights as probability for traversal, and calculating the probability of the random walker arriving at each node. This probability is then taken as a node's centrality.

HITS [Kle99] calculates two complementary metrics for each node in the network: an *authority* score, which quantifies how much each node serves as an authority for referencing other nodes; and a *hub* score, which quantifies how well a node refers to authorities. We will use both hub and authority scores for this evaluation.

5.2.4 Dependent variables

To investigate how the choice of network and centrality metric affect ranking, we will measure the normalized discounted cumulative gain (nDCG), a commonly used metric for evaluating the performance of ranking. We will summarize the nDCG by taking the mean across datasets, such that we can compare the choice of network and centrality.

5.2.5 Results

In Table 5.1, we show the mean nDCG for each network type and centrality metric, per dataset. We observe that the best combination of network and metric varies significantly per dataset. However, in four of the five datasets (aerospace, compsci, datascience, and medicine) some centrality metric applied to the interaction network obtains the highest nDCG score.

When we look at the summary of the results across datasets, as displayed in Table 5.2, a clearer picture emerges. For each centrality metric, the mean nDCG score is significantly higher when applied to interaction networks instead of co-occurrence networks. Here, we see that eigenvector centrality and PageRank perform best across datasets, closely followed by HITS authorities and Katz centrality.

A visualization of the resultant interaction networks, highlighting the five most central nodes per dataset as obtained through PageRank, can be found in Figure 5.1.

5.3 Discussion

We constructed a meme interaction network and a meme co-occurrence network for the top 100 memes identified through our Memeseeker method, and subsequently ranked these memes by their centrality scores as measured by six distinct centrality metrics. Through Tables 5.1 and 5.2, we find support for the correctness of our method for the identification of meme interactions, through validation of our hypothesis that centrality in meme interaction networks better quantifies perceived relevance than in co-occurrence networks, for most centrality measures, on most datasets.

Through comparison of Table 5.2 with Table 4.5 (summarizing the results of the online evaluation of meme identification), we further find support that after meme identification, re-ranking of memes by the centrality in their interaction networks allows for somewhat better extraction of key concepts from online communities: ranking by propagation score obtains a mean nDCG of ca 0.7, whereas ranking by e.g. PageRank centrality obtains a mean nDCG of ca 0.8.

community	metric	interaction network	co-occurrence network
aerospace	betweenness	0.795	0.836
aerospace	eigenvector	**0.869**	0.849
aerospace	katz	0.827	**0.863**
aerospace	pagerank	0.859	**0.868**
aerospace	hits-hubs	0.834	0.849
aerospace	hits-authorities	0.830	0.849
compsci	betweenness	**0.796**	0.739
compsci	eigenvector	**0.783**	0.676
compsci	katz	**0.782**	0.665
compsci	pagerank	0.778	0.748
compsci	hits-hubs	0.739	0.682
compsci	hits-authorities	0.780	0.686
datascience	betweenness	0.764	0.790
datascience	eigenvector	**0.872**	0.800
datascience	katz	0.870	0.809
datascience	pagerank	**0.878**	0.811
datascience	hits-hubs	0.805	0.800
datascience	hits-authorities	**0.872**	0.800
medicine	betweenness	0.785	0.795
medicine	eigenvector	**0.823**	0.758
medicine	katz	**0.826**	0.763
medicine	pagerank	0.822	0.761
medicine	hits-hubs	0.790	0.758
medicine	hits-authorities	**0.827**	0.758
psychology	betweenness	0.676	0.634
psychology	eigenvector	0.669	**0.684**
psychology	katz	0.676	0.644
psychology	pagerank	0.680	0.606
psychology	hits-hubs	0.645	**0.684**
psychology	hits-authorities	0.677	**0.684**

Tab. 5.1.: Results of the evaluation. This table shows the mean nDCG scores for both types of networks and each centrality metric, per dataset. Here, the three best performing ranking methods are highlighted per dataset.

metric	interaction network	co-occurrence network
betweenness	0.763	0.759
eigenvector	**0.803**	0.754
katz	0.796	0.749
pagerank	**0.803**	0.759
hits-hubs	0.763	0.755
hits-authorities	**0.797**	0.756

Tab. 5.2.: Summary of the evaluation. This table shows the mean of mean nDCG scores across datasets, for both types of networks and each centrality metric. Here, the best three performing ranking methods are highlighted.

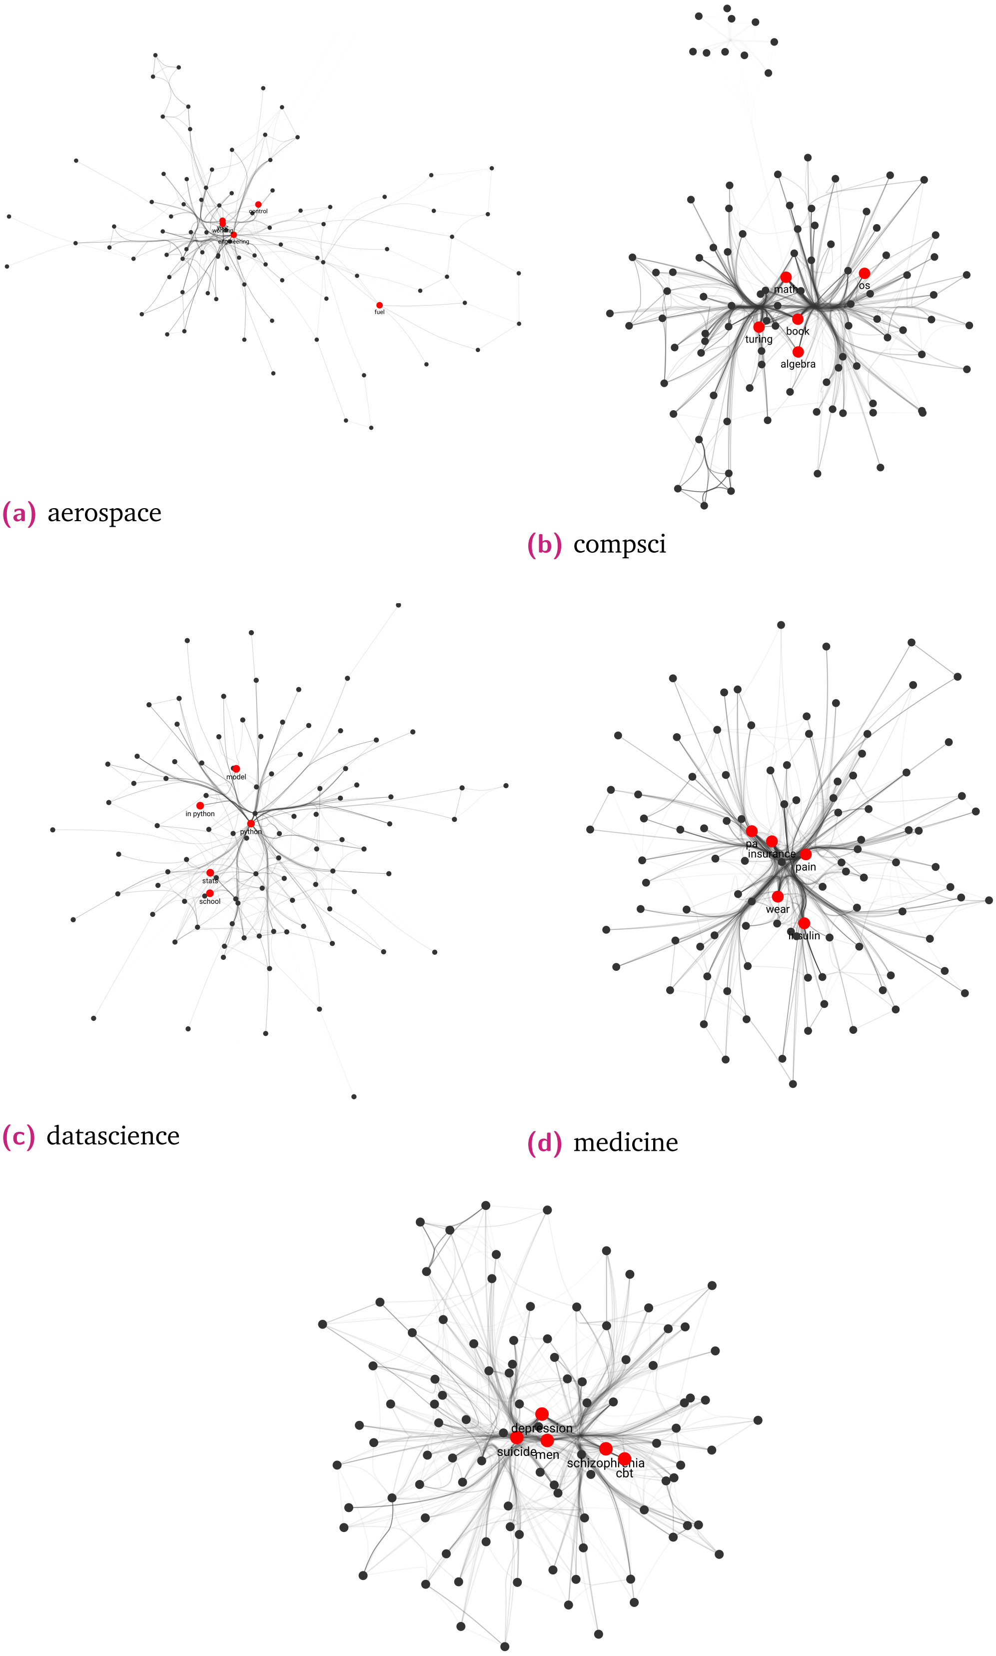

(a) aerospace

(b) compsci

(c) datascience

(d) medicine

(e) psychology

Fig. 5.1.: Meme interaction networks for each dataset. Here, for each dataset, the five most central memes as measured through PageRank are displayed in red. As the original interaction networks are very dense, which makes the visualization difficult, we have pruned the networks using the disparity filtering method [Ser+09].

6 Conclusion

> *There is a struggle for existence because a vast array of memes is competing for the limited resource of human attention, and therefore the fitness of any given meme will be influenced chiefly by its ability to gain and retain attention.*
>
> — **Kate Distin**
> Author of 'The Selfish Meme'

We constructed a method for the large scale identification of scientific memes that propagate through science-oriented online communities, and showed how this method may be slightly altered in order to not only identify memes themselves, but also the interactions in which they engage. We further showed that memes propagating through such online communities take the form of concepts which are relevant to the communities at hand. In this Chapter, we will return to the research questions as we posed them at the beginning of Chapter 1, and address them individually in light of proposed theoretical contributions and the conducted evaluations.

6.1 RQ1: Can we identify memes propagating through online communities?

In Chapter 1, we set out to construct a method for the large-scale identification of memes which propagate through online communities. We argued that such a method would be a stepping stone for contemporary scientific endeavor in the field of memetics. To address RQ1, in Section 3.2, we introduced the Memeseeker method for the identification of these memes, consisting of three steps: construction of a *propagation network* through which memes could propagate; assigning each pattern a *propagation score*, quantifying how much of its propagation could be explained by imitation; and *selection* of patterns as memes for which this propagation score is significant. We conducted a coarse, large-scale offline evaluation by applying this method to Reddit and StackExchange data, and an finer, small-scale online evaluation on Reddit data using human participants.

Through the offline evaluations conducted in Section 4.2 and Appendix C, we found that the proposed Memeseeker method is able to extract *relevant* ngrams — corresponding to Wikipedia titles and not being stopwords — from Reddit and StackExchange datasets with a reasonably high precision, with relevant ngrams obtaining higher propagation scores than irrelevant ones (c.f., Figures 4.2 and C.2, respectively). The online evaluation conducted in Section 4.3 found further support that the proposed method is able to extract concepts relevant to particular scientific fields. The proposed method scales to web scale datasets, with the evaluations being run on datasets containing tens to hundreds of thousands distinct users and posts, selecting thousands of memes from tens of millions of candidate behavioral patterns in under half an hour.

With regards to the calculation of the propagation score, we have compared a number of distinct alternatives. Our method requires the measurement of meme-carrying *propagation vehicles*, for which memetics predicts *egos* and *artefacts*. Furthermore, we can quantify the effect that influence of a meme has on the probability of carrying that meme through the *absolute risk*, *relative risk*, and *odds ratio*. We found that measurement of egos using the absolute risk consistently performed best with respect to the identification of relevant concepts to scientific fields.

Under the assumption that memes propagating through science-oriented communities indeed take the form of concepts relevant to particular scientific fields, we can address RQ1: we have demonstrated that we were able to identify memes propagating through online communities on a population scale.

6.2 RQ2: Can we identify the interactions between these memes?

To address RQ2, in Section 3.3, we proposed an alteration of the Memeseeker method to quantify the interactions between memes, by quantifying the effect that influence of some meme has on the probability of carrying some *other* meme. However, the identification of meme interactions is a novel task, making direct evaluation of the correctness of this method difficult.

We instead analyzed properties of the *meme interaction networks* by way of their similarity to other types of interaction networks. We evaluated whether *centrality* of memes within their interaction networks is an indicator for perceived relevance, through a comparison with co-occurrence networks. We found that a meme's centrality in meme interaction networks is indeed a better indicator of perceived relevance compared to its centrality in meme co-occurrence networks, for all tested centrality metrics. This allows us to address RQ2 with some confidence: we found support that our method identifies the interactions between memes.

6.3 RQ3: Can we leverage memetic theory to extract key concepts from online communities?

Through the evaluation in Chapter 4 we have shown that our method is able to identify relevant concepts to particular scientific fields, and rank them relatively well according to perceived relevance. Through our evaluation of the identification of meme interactions, we found we can significantly improve upon this ranking, by leveraging the centrality of memes within their meme interaction networks, with eigenvector centrality and PageRank performing best for this task. The relatively precise identification of our meme identification method, combined with the improved method for ranking by meme centrality, allows us to address RQ3: we have demonstrated a novel way for the extraction of key concepts from online communities.

6.4 Limitation and future work

In this thesis we made a number of simplifying assumptions regarding exposure. Firstly, a user leaving a comment on a thread was assumed to be exposed to all comments in that thread. Secondly, we applied a close-world assumption; that exposure solely occurred through the Reddit platform, i.e. external exposure was not considered. Thirdly, we only considered binary expression and exposure. In epidemiology, dose-response analysis allows for finer-grained investigation of the effect of exposure on expression. Furthermore, it has been suggested that behavior spreads through *complex contagion* [Rom+11], where both the amount and the source of exposure play a role in propagation. In future work, these principles should be taken into account, which might allow the identification of memes with higher accuracy.

Regarding the evaluations, the automated evaluations in Section 4.2 and Appendix C used a relatively broad measure of relevance, and the human evaluations in Sections 4.3 and 5.2 had a relatively small number of participants. Nonetheless, the evaluations do support each other in terms of results, making up for some of the broadness and lack of data.

Furthermore, our study is limited to the culture and correspondence structure inherent to Reddit and StackExchange. We do expect our method to extend well to other types of data, as long as we can determine expression of and exposure to traceable behavioral patterns.

We can now identify memes, and the interactions they engage in, as they propagate through online communities, on a population scale and with a reasonably high precision. This finding allows for further validation of memetics' theories. We further hope that this allows the theory to contribute to contemporary scientific endeavor in other fields, such as information extraction and social network analysis. One such contribution could be the study of memeplexes through clustering of meme interaction networks, which may provide alternatives for topic modeling techniques such as LSI or LDA.

Another interesting direction is to distill meme pathways from meme interaction networks. As was discussed in Section 2.5, pathway analysis is commonly used to find protein pathways in protein interaction networks, which sheds light on the sequential processing of metabolites within cells. Such techniques could also be applied to meme interaction networks, where they may shed light on common "learning trajectories" within communities of individuals. We hypothesize that such learning trajectories take the form of sequences of memes, or even memeplexes, that are often acquired in order. It may be interesting to contrast these "natural" learning trajectories with the curricula provided by formal education, to see where they may differ or overlap.

Finally, our current work could contribute to social network analysis, by helping to shed light on online communities themselves. Through laying bare the propagation pathways over which memes commonly propagate, this may allow us to find which users on these platforms are introducing new memes, which users are important conduits in this spread, and which users only imitate what they see.

Symbols and derivations

A

A.1 Basics

Memes

$M = \{m_1, m_2, ..., m_{|M|}\}$

Artefacts

$A = \{a_1, a_2, ..., a_{|A|}\}; a_1, a_2, ..., a_{|A|} \subseteq M$

Egos

$E = \{e_1, e_2, ..., e_{|E|}\}; e_1, e_2, ..., e_{|E|} \subseteq A$

A.2 Exposure and expression

Exposure to some artefact of some ego, at some real-valued point in time

$X = \{x_1, x_2, ..., x_{|X|}\}; x_1, x_2, ..., x_{|X|} \in A \times E \times \mathbb{R}$

Expression of some artefact by some ego, at some real-valued point in time

$Y = \{y_1, y_2, ..., y_{|Y|}\}; y_1, y_2, ..., y_{|Y|} \in A \times E \times \mathbb{R}$

Artefact involved in some exposure or expression

$z^{(a)} \in A; z \in X \cup Y$

Ego involved in some exposure or expression

$z^{(e)} \in E; z \in X \cup Y$

Time at which some exposure or expression occurred

$z^{(t)} \in \mathbb{R}; z \in X \cup Y$

A.3 Influence

Influence by one artefact on another

$$a_1 \rightarrow a_2 \Leftrightarrow (\exists x \in X; \exists y \in Y)[x^{(a)} = a_1 \wedge y^{(a)} = a_2 \wedge x^{(e)} = y^{(e)} \wedge x^{(t)} < y^{(t)}]$$

A.4 Artefact sub-populations

Artefacts which were influenced by meme *i* and carry meme *j*

$$A_j^i = \{a \in A : j \in a \wedge (\exists a' \in A)[i \in a' \wedge a' \rightarrow a]\}; i, j \in M$$

Artefacts which were not influenced by meme *i* and carry meme *j*

$$A_j^{\not{i}} = \{a \in A : j \in a \wedge (\neg\exists a' \in A)[i \in a' \wedge a' \rightarrow a]\}; i, j \in M$$

Artefacts which were influenced by meme *i* and do not carry meme *j*

$$A_{\not{j}}^i = \{a \in A : j \notin a \wedge (\exists a' \in A)[i \in a' \wedge a' \rightarrow a]\}; i, j \in M$$

Artefacts which were not influenced by meme *i* and do not carry meme *j*

$$A_{\not{j}}^{\not{i}} = \{a \notin A : j \in a \wedge (\neg\exists a' \in A)[i \in a' \wedge a' \rightarrow a]\}; i, j \in M$$

A.5 Ego sub-populations

Egos which were influenced by meme *i* and carry meme *j*

$$E_j^i = \{e \in E : (\exists a \in e)[a \in A_j^i] \wedge (\neg\exists a' \in e)[a' \in A_j^{\not{i}}]\}; i, j \in M$$

Egos which were not influenced by meme *i* and carry meme *j*

$$E_j^{\not{i}} = \{e \in E : (\exists a \in e)[a \in A_j^{\not{i}}]\}; i, j \in M$$

Egos which were influenced by meme *i* and do not carry meme *j*

$$E_{\not{j}}^i = \{e \in E : (\exists a \in e)[a \in A_{\not{j}}^i] \wedge (\neg\exists a' \in e)[a' \in A_j^i \vee a' \in A_j^{\not{i}}]\}; i, j \in M$$

Egos which were not influenced by meme *i* and do not carry meme *j*

$$E_{\not{j}}^{\not{i}} = \{e \in E : (\exists a \in e)[a \in A_{\not{j}}^{\not{i}}] \wedge (\neg\exists a' \in e)[a' \in A_j^i \vee a' \in A_j^{\not{i}}] \vee a' \in A_{\not{j}}^i\}; i, j \in M$$

B Selected memes from Reddit datasets

rank	aerospace	compsci	datascience	medicine	psychology
1	catia	lush	python 2	dpc	peep show
2	delft	bullshit generator markov	vim	inh	peep
3	fortran	reddit bullshit	galvanize	varus	portal
4	itar	generator markov chains	phd	23andme	ubi
5	blades	reddit bullshit generator	aws	ketamine	bitcoin
6	citizenship	has chance of being	data table	io	tm
7	matlab	has chance of	tableau	cprs	mdma
8	python	bullshit generator markov chains	mba	cvp	queer
9	combustion	reddit bullshit generator markov	bi	egdt	emdr
10	security clearance	has chance	excel	ai	hate speech
11	supersonic	generator markov	dplyr	dengue	selfies
12	china	sensible to	d3	circumcision	free will
13	airfoil	xen	clusters	chaperone	microaggressions
14	book	free will	pandas	scribes	spanking
15	clearance	pumping	r	spider	tsa
16	stealth	exptime	book	iud	bpd
17	the shuttle	nobel	etl	ferritin	conversion therapy
18	cad	pumping lemma	rstudio	tpa	of free will
19	fuel	utm	memory	ppd	aba
20	gravity	women	cs	omfs	stereotype threat
21	3d	cache misses	sas	crp	iq
22	flow	np	kaggle	pslf	free will is

rank	aerospace	compsci	datascience	medicine	psychology
23	cfd	gc	github	ibr	meditation
24	mach	aws	of living	ebola	rorschach
25	security	poker	python	sirs	cbt
26	calculus	fpga	europe	lr	adhd
27	grad	financial aid	vba	cjd	mbti
28	pressure	lemma	algebra	pots	hypnosis
29	thrust	md5	internships	cerner	nicotine
30	math	the pumping lemma	internship	iuds	blank slate
31	velocity	trie	sql	guns	eq
32	c++	the pumping	stats	bicarb	a psyd
33	georgia tech	jit	windows	bubbles	wellbutrin
34	georgia	hash	the model	breastfeeding	ketamine
35	propulsion	rust	random forest	zika	tattoos
36	mars	bloom	column	gun	gre
37	physics	quantum	anaconda	fgm	will is
38	degree	ternary	pca	single payer	bot
39	language	kay	program	sperm	conspiracy
40	spacex	turing	matlab	esr	big 5
41	weight	hash function	columns	ross	p-value
42	heat	genetic	masters	toradol	spss
43	grad school	password	interview	cbt	nlp
44	phd	gpu	regression	scrubs	the gre
45	school	latex	js	abortion	placebo
46	drag	git	spark	coats	learning styles
47	an internship	halting	java	lupus	pen
48	mass	erlang	scala	vaccine	iq tests
49	efficiency	halting problem	r is	suboxone	homeopathy
50	wing	regex	the phd	omm	vegetarian
51	payload	xml	math	cfs	behaviorism
52	earth	polynomial	forest	expiration	amendment
53	lift	dfa	bias	torsades	feminist
54	stability	annealing	healthcare	bariatric	lucid
55	purdue	cuda	git	procalcitonin	introvert
56	resume	neurons	score	peep	epigenetics
57	nasa	nsa	degree	lasik	dbt
58	citizen	p	hadoop	compressions	depression
59	ae	watson	udemy	android	racism
60	the wing	chess	scraping	vaccines	mindfulness
61	internship	monads	api	memes	transgender
62	schools	wolfram	courses	adenosine	i/o
63	f-35	turing complete	course	epic is	atheism
64	speed	cobol	variables	pharmacists	quantum

rank	aerospace	compsci	datascience	medicine	psychology
65	article	algebra	linux	va	synesthesia
66	pilots	ethical	linear algebra	peritonitis	morality is
67	engines	strong ai	pm	d-dimer	slate
68	companies	in np	article	insurance	emergent
69	learn	hex	a phd	keto	heuristics
70	industry	tex	bootcamp	hipaa	apple
71	classes	consciousness	ms	historian	violent video
72	gpa	entropy	trees	allergy	exposure therapy
73	rocket	dragon book	ai	inhalers	porn
74	landing	calc	calculus	em	circumcision
75	energy	the halting	ml	payer	gun control
76	center	dragon	the program	stethoscope	aa
77	aero	the halting problem	customers	maggots	masculinity
78	pilot	compression	spss	a chaperone	conversion
79	engineering	book	ram	picc	conservatives
80	shuttle	ads	function	f1	npd
81	low	quicksort	rows	hpv	replication
82	career	mips	shiny	pa	freud
83	plane	lambda	cluster	aed	psyd
84	college	heuristic	linear	epic	suicide is
85	told	math	bayesian	software	psychoanalysis
86	experience	lambda calculus	jupyter	bipap	schizophrenia
87	control	haskell	cv	pagers	the mbti
88	year	recursion	linkedin	pain	ptsd
89	team	tail	values	wear	men
90	masters	syntactic	school	narcan	games
91	wings	quantum computer	in python	propofol	the placebo
92	why	os	model	honey	addiction
93	project	turing machine	classification	modafinil	gaming
94	program	isomorphism	resume	concierge	lsd
95	bad	nlp	features	bls	suicide
96	advice	sipser	insight	epi	violent video games
97	interview	linux	manager	acupuncture	feminism
98	a rocket	linear algebra	academia	nnt	gay
99	working	np-complete	c++	imgs	revenge
100	systems	o(1)	logistic	insulin	racist

Tab. B.1.: Top 100 memes from all fields, extracted using the Memeseeker method with egos as vehicles and absolute risk as effect, as discussed in the evaluation in Chapter 4.

C Evaluation of meme identification on StackExchange datasets

In this Appendix, we will replicate the offline evaluation as conducted in Section 4.2 on data extracted from five science-oriented communities on the online questions and answers platform StackExchange. The purpose of this replication is to validate that the Memeseeker method not only works on Reddit, but across multiple platforms. The indepedent variables, dependent variables, and hypotheses will not be repeated in this evaluation.

C.1 Procedure

Again, we evaluate our method in a similar fashion as information retrieval problems, where we rank the selected memes by their propagation score, and use relevancy-based metrics to quantify the performance of the system. We use the same automated procedure to assign relevancy to each selected meme as was discussed in 4.2: we say that an ngram is *relevant* when it is a the title of a Wikipedia article, and not a stopword — as judged by the Python Natural Language Toolkit [Bir+09].

Dataset We acquire five datasets, each containing all questions, answers, and comments created within a specific science-oriented community of StackExchange: chemistry (containing discussions around Chemistry); cs (Computer Science); datascience (Data Science); mathoverflow (Mathematics); stats (Statistics). We obtain these datasets from a publicly available dump of StackExchange data on Archive.org. See Table C.1 for a small summary of the datasets.

Text cleaning We perform a simple cleaning process on the textual contents of the comments. We remove double quotes ("), periods (.), commas (,), exclamation marks (!) and question marks (?). All text is converted into lowercase.

Egos, artefacts and behavioral patterns Per dataset, each question, answer, or comment represents a single artefact, which was expressed by an ego — the user who created it. The behavioral patterns from which we will be selecting and scoring memes are all unique ngrams of length 1 through 4, extracted from the textual contents of the comments. A summary of the data, displaying the number of egos, artefacts, and unique ngrams per dataset can be found in Table C.1.

community	# egos	# artefacts	# ngrams
chemistry	15,709	142,083	16,556,741
cs	17,616	152,350	18,311,709
datascience	6,469	34,357	5,363,771
mathoverflow	38,013	704,743	72,954,919
stats	69,862	659,805	74,019,990

Tab. C.1.: Summary of datasets, showing the number of egos, artefacts and unique ngrams for each dataset.

Constructing the propagation network In order to construct a propagation network for our StackExchange datasets, we must determine which ngrams an ego has expressed, and to which ngrams an ego was exposed. Establishing expression is trivial, as each artefact expresses the ngrams it contains. However, as the datasets do not contain any information on which users have read which comments, establishing exposure is less trivial, and requires us to develop a heuristic.

On StackExchange, users can create new questions related the community's field, to which other users can respond with an answer. For both questions and answers, there is a comment thread in which users can ask for or provide clarification. This correspondence by users on StackEchange thus forms a hierarchy which is structured like a tree, as displayed in Figure C.1. An artefact's ancestors are those artefacts which are higher up this hierarchy. We heuristically assume that a user creating an artefact was exposed to all ancestors of that artefact. As such, we can construct a propagation network by connecting each user's comments to their ancestors, and to the ancestors of artefacts that were created earlier by that user.

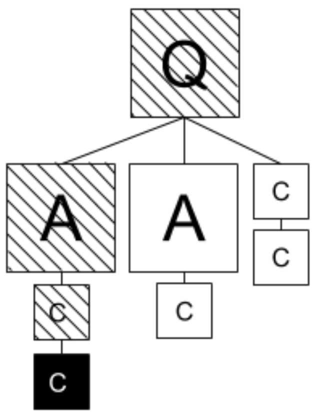

Fig. C.1.: Example of the hierarchy of correspondence on StackExchange. In this image, squares containing a Q represent questions; A represent answers; and C represent comments. We heuristically assume that a user is exposed to all ancestors of artefacts they create. In this image, the user creating the comment marked in black is assumed to have been exposed to the shaded artefacts.

C.2 Results

We compare the results for vehicle (ego or artefact) and effect (absolute risk, relative risk, or odds ratio); resulting in 6 conditions. Comparisons are given for five datasets (chemistry, cs, datascience, mathoverflow, stats). The results of the evaluation are summarized per dataset in Table C.2, and a summary across datasets is given in Table C.3. A small manual selection of memes, as ranked and selected through ego vehicles and absolute risk, can be found in Table C.4. For the top 100 memes, as ranked and selected through ego vehicles and absolute risk, we refer the reader to Appendix D.

community	vehicle	effect	# memes (N)	nDCG @100	nDCG @N	precision @100	precision @N
chemistry	egos	AR	9,306	0.822	0.950	0.870	0.438
chemistry	egos	RR	9,306	0.752	0.926	0.780	0.438
chemistry	egos	OR	9,306	0.770	0.930	0.780	0.438
chemistry	artefacts	AR	6,411	0.643	0.894	0.670	0.430
chemistry	artefacts	RR	6,411	0.000	0.820	0.000	0.430
chemistry	artefacts	OR	6,411	0.000	0.820	0.000	0.430
cs	egos	AR	7,437	0.797	0.937	0.850	0.422
cs	egos	RR	7,437	0.740	0.915	0.780	0.422
cs	egos	OR	7,437	0.729	0.916	0.780	0.422
cs	artefacts	AR	7,096	0.333	0.830	0.420	0.242
cs	artefacts	RR	7,096	0.000	0.730	0.000	0.242
cs	artefacts	OR	7,096	0.000	0.730	0.000	0.242
datascience	egos	AR	1,064	0.893	0.972	0.930	0.745
datascience	egos	RR	1,064	0.812	0.959	0.860	0.745
datascience	egos	OR	1,064	0.837	0.965	0.860	0.745
datascience	artefacts	AR	668	0.676	0.914	0.720	0.669
datascience	artefacts	RR	668	0.350	0.861	0.430	0.669
datascience	artefacts	OR	668	0.350	0.861	0.430	0.669
mathoverflow	egos	AR	53,242	0.333	0.900	0.410	0.197
mathoverflow	egos	RR	53,242	0.387	0.863	0.480	0.197
mathoverflow	egos	OR	53,242	0.369	0.866	0.460	0.197
mathoverflow	artefacts	AR	46,617	0.499	0.896	0.570	0.251
mathoverflow	artefacts	RR	46,617	0.023	0.844	0.030	0.251
mathoverflow	artefacts	OR	46,617	0.023	0.844	0.030	0.251
stats	egos	AR	32,908	0.678	0.924	0.750	0.241
stats	egos	RR	32,908	0.508	0.875	0.580	0.241
stats	egos	OR	32,908	0.515	0.877	0.580	0.241
stats	artefacts	AR	27,903	0.366	0.879	0.390	0.227
stats	artefacts	RR	27,903	0.030	0.792	0.010	0.227
stats	artefacts	OR	27,903	0.048	0.793	0.010	0.227

Tab. C.2.: Results of the offline evaluation. This table shows the nDCG and precision, evaluated at rank 100 and at the final rank N (the total number of selected memes) for each combination of vehicle and effect, per dataset.

Ranking and selection perform best for egos, compared to artefacts We observe that measurement of egos again significantly outperforms measurement of artefacts, both in terms of nDCG and precision. This difference is especially pronounced for measurement of relative risk and odds ratio. In Figure C.2 we show the progression of precision over rank,

vehicle	effect	mean nDCG @100	mean nDCG @N	mean precision @100	mean precision @N
egos	AR	0.705	0.936	0.762	0.409
egos	RR	0.640	0.908	0.696	0.409
egos	OR	0.644	0.911	0.692	0.409
artefacts	AR	0.503	0.883	0.554	0.364
artefacts	RR	0.081	0.810	0.094	0.364
artefacts	OR	0.084	0.810	0.094	0.364

Tab. C.3.: Summary of the offline evaluation. This table shows the mean nDCG and mean precision across datasets, evaluated at rank 100 and at the final rank N, for each combination of vehicle and effect.

community	rank	propagation score	meme
chemistry	13	0.540	inert pair
chemistry	26	0.481	the triple point
chemistry	35	0.461	ice
cs	18	0.598	semaphore
cs	28	0.499	md5
cs	36	0.480	morse code
datascience	6	0.462	churn
datascience	14	0.395	lda
datascience	43	0.327	feature selection
mathoverflow	8	0.999	sub-laplacian
mathoverflow	23	0.833	base bumping
mathoverflow	50	0.526	quantum channels
stats	4	0.875	nonconforming
stats	21	0.463	modularity
stats	53	0.366	reward function

Tab. C.4.: A manual selection of memes for each dataset, for illustrative purposes. These memes were selected from the top 100 memes as ranked and selected through ego measurement and absolute risk. The entire list can be found in Appendix D.

which shows low precision of artefact scores at low ranks, only to increase drastically for higher ranks. This result is also visible in Table C.2 and Table C.3, where we observe that precision@100 values are extremely low for artefacts through measurement of relative risk and odds ratio, and somewhat low for measurement of absolute risk.

Ranking performs best for absolute risk, compared to relative risk and odds ratio

For both ego and artefact measurement, absolute risk outperforms relative risk and odds ratio in terms of nDCG, and relative risk and odds ratio appear to perform similar.

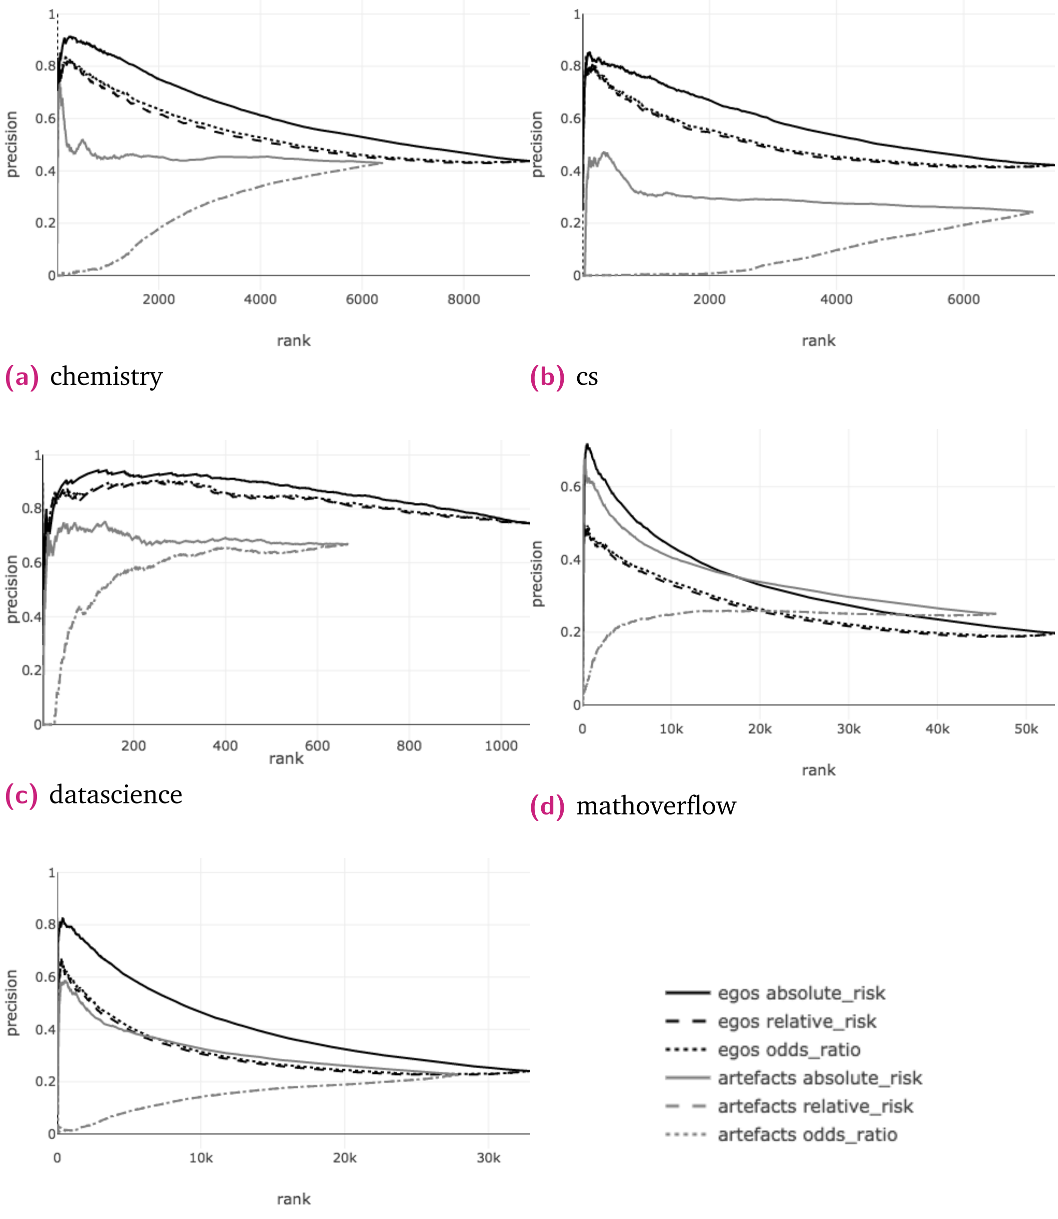

(a) chemistry (b) cs

(c) datascience (d) mathoverflow

(e) stats

Fig. C.2.: Offline evaluation: rank vs. precision for each combination of vehicle and effect, per dataset. Here, black lines denote ego vehicles, grey lines denote artefact vehicles. Solid lines denote absolute risk, dashed lines denote relative risk, and dotted lines denote odds ratio. As relative risk and odds ratio perform very similar, they overlap in these figures.

C.3 Discussion

This replication of the offline evaluation on StackExchange datasets yields similar results as the original study on Reddit datasets as performed in Section 4.2, in that we find measurement of ego propagation outperforming artefact propagation, and absolute risk outperforming relative risk and the odds ratio, both in terms of nDCG and precision (c.f., Table C.3). However, in contrast to the evaluation on Reddit data, we find through Table C.1 that the quantification of artefact propagation through the absolute risk achieves a performance close to ego propagation, especially for lower ranks. This result can also be clearly seen in Figure C.2, which shows a clear difference with Figure 4.2 of the Reddit evaluation, where the quantification of artefact propagation through absolute risk does not significantly outperform other effects. This may be due to StackExchange not being host to so many bots as Reddit, skewing the results less in the favor of ego vehicles. This further lends support for the hypothesis stated in 4.2 that absolute risk should outperform relative risk and the odds ratio.

Another difference we notice between the evaluation on Reddit data in Section 4.2 and the current evaluation is that ego propagation selects *fewer* memes than artefact propagation for Reddit data, whereas it selects *more* memes for StackExchange data (c.f., Table 4.2 and Table C.2, respectively). This result, too, may be caused by the relative lack of bots on StackExchange, or it may be due to different correspondence structure (i.e. general discussion on Reddit versus questions and answers on StackExchange), or different discussion norms (i.e. more informal on Reddit). Future research could shed more light on this phenomenon.

D

Selected memes from StackExchange datasets

rank	chemistry	cs	datascience	mathoverflow	stats
1	leucoindigo	automatic variable	graph embedding	$k^{p	the star schema
2	chaperones	unipathic	string3	$k^{p q}(x)$	|| [3]
3	congo red	permutation rules	string2	potential diagonalizability	$so(3)$
4	hepes	stack inspection	brat	the charney-davis conjecture	nonconforming
5	congo	diploid	hellinger distance	hubbiness	the joker
6	tbaf	initial temperature	churn	charney-davis conjecture	entity resolution
7	chloroacetic	unproductive	gini	axiomatic rank	opting in
8	molar absorptivity	cosine similarity	rdd	sub-laplacian	@gbow28
9	malonic	than human	cosine	q^h)$	@wildetudor
10	the nodal	sentential forms of	survival	the charney-davis	hausdorff distance
11	state approximation	convoy effect	smote	banach integral	kerrich
12	zwitterion	external fragmentation	word2vec	q^h)$ is	lcs
13	inert pair	wedge	with missing	almost malnormal	the reward function
14	boric acid	pca	lda	cohesive set	smc
15	steam distillation	scale-free	spark	x^{**})$	frog
16	ascorbic acid	deadlock	ner	basis constant	benford's law
17	ascorbic	reversible gates	learning rate	zombies	benford's
18	boric	semaphore	cosine similarity	popovici	nps
19	milk	rpn	pca	nyldon	frogs
20	fortran	mapreduce	gpu	nyldon words	icc
21	aspartame	grover's	convolution	morse number	modularity
22	methyl orange	a dpda	batch	higgs mechanism	balls

rank	chemistry	cs	datascience	mathoverflow	stats
23	caffeine	phone number	twitter	base-bumping	mediation
24	benzyne	eve	relu	ultrafield	pca
25	hypervalency	convex hull	reinforcement	an ultrafield	marbles
26	the triple point	hull	documents	minimally unprovable	julia
27	inert pair effect	cisc	clustering	the evader	sankey
28	of fusion	md5	missing values	centro-affine	nmf
29	furan	utm	similarity	$[k^2	the som
30	nitride	starvation	svm	evader	a martingale
31	carbocation	pumping	weather	centro-affine curvature	cca
32	chemisorption	sticks	distance	$k[x]^g$	t-sne
33	graphene	generating function	excel	approximately inner	sweave
34	eyring	mealy	overfitting	ilmanen	hoeffding's
35	ice	broadcast	svd	robber	dtw
36	thermite	morse code	lstm	olives	lstm
37	trans	morse	sentiment	legal subsets	rao-blackwell
38	pair effect	sudoku	seasonality	of the fibonacci word	conjoint
39	gc	polygon	spam	monadic functors	lasso
40	entropy	sentential	time series	zipf's	roc
41	dipole moment	logical address	ip	zipf's law	manova
42	blood	perceptron	rnn	segre classes	earthquakes
43	dipole	nfa	feature selection	screensaver	elasticity
44	molality	pumping lemma	anomaly	analytic semigroup	mds
45	tertiary	dfa	correlation	thief	raters
46	mercury	decompression	orange	eigencurve	ump
47	buffer	pda	reinforcement learning	the moore method	marble
48	wurtz	page size	k-means	ahss	viterbi
49	xrd	tsp	xgboost	high-volume	periodogram
50	triple point	ll(1)	sgd	quantum channels	cards
51	azide	stone	tree	the thief	clustering
52	phenol	overfitting	ratings	willmore	imputation
53	chiral	page table	feature engineering	stationary sets	reward function
54	methanol	category theory	categorical	@sergeiakbarov	of balls
55	calorimeter	mst	the correlation	co-hopfian	the icc
56	pyridine	the convex	cluster	mereology	lda
57	salt bridge	diameter	tf-idf	tarski monster	vecm

rank	chemistry	cs	datascience	mathoverflow	stats
58	grignard	ctl	outliers	the busemann	vif
59	ph	the convex hull	document	diffeological spaces	mahalanobis
60	wine	rectangles	series	cosheaf	baum-welch
61	oxidation state	pcp	df	tutte polynomial	2sls
62	benzoic	spanning	caret	copula	poisson
63	the calorimeter	residual	filters	deduction theorem	auc
64	tlc	dpda	cnn	tree property	glmnet
65	resonance	dcfl	images	the tutte polynomial	survival
66	ortho	quantum	the learning rate	relatively hyperbolic	the whiskers
67	anthracene	clique	gradient	simple homotopy equivalence	vc dimension
68	citric	mutual exclusion	word	the mahler measure	bonferroni
69	cyclobutadiene	context switch	pooling	ibn	clusters
70	benzoic acid	subtyping	rating	curvature flow	data augmentation
71	pyrite	von neumann	words	vixra	a pie chart
72	acetone	the cache is	rate	a perverse sheaf	relu
73	yeast	turing test	confusion matrix	spectral sequence	the propensity score
74	thionyl	steiner	sentiment analysis	hopfian	amos
75	distillation	interpolation	weights	cofibrations	mutual information
76	quartet	trie	outlier	prikry	@aksakal
77	the epoxide	risc	layer	skein relation	gee
78	cis	chess	precision	brownian bridge	fdr
79	baking	euler	text	hilbert transform	forward-backward
80	radius	fragmentation	selection	fibrant	price elasticity
81	gallium	hash	install	bbd	bmi
82	refractive	lemma	logistic regression	moore method	sur
83	enol	bipartite	clusters	the ahss	gretl
84	real gas	convex	hour	the bergman	the ljung-box
85	diastereomers	cfgs	big data	perfectoid	whiskers
86	balloon	grammar	anomaly detection	nno	gini
87	hydride shift	lba	matrix	virtually free	urn

rank	chemistry	cs	datascience	mathoverflow	stats
88	edta	regular grammar	hadoop	the steenrod	garch
89	wurtz reaction	the pumping	feature extraction	t-structures	ridge
90	thionyl chloride	spanning tree	detection	operad	aic
91	ester	comparator	validation	the radon transform	pie chart
92	benzene	exptime	entropy	contact structure	tensors
93	dichloromethane	halting	regression	model structure	cluster
94	egg	independent set	reduction	helly's theorem	martingale
95	uranium	polygons	softmax	marginals	@stéphanelau-rent
96	aromaticity	hamiltonian	sigmoid	abelian scheme	mcar
97	vinegar	coins	unsupervised	r-matrix	ljung-box
98	covalent	interrupts	plot	cofibrant	jags
99	chromate	deep learning	sentence	spanning trees	beta regression
100	nucleophile	convolution	csv	bundle	k-means

Tab. D.1.: Top 100 memes from all fields, extracted using the Memeseeker method with egos as vehicles and absolute risk as effect, as discussed in the evaluation in Appendix C.

Bibliography

[AC00] Alan Agresti and Brian Caffo. „Simple and effective confidence intervals for proportions and differences of proportions result from adding two successes and two failures“. In: *The American Statistician* 54.4 (2000), pp. 280–288 (cit. on p. 17).

[Agr07] Alan Agresti. *An introduction to categorical data analysis, Second Edition*. Wiley New York, 2007. Chap. 2 (cit. on pp. 15, 17, 18).

[Bes97] Michael L Best. „Models for interacting populations of memes: Competition and niche behavior“. In: *The Fourth European Conference on Artificial Life. The MIT Press, Cambridge MA*. 1997 (cit. on pp. 8, 17).

[BF+17] Héctor Beck-Fernandez, David F Nettleton, Lorena Recalde, Diego Saez-Trumper, and Alexis Barahona-Peñaranda. „A system for extracting and comparing memes in online forums“. In: *Expert Systems with Applications* 82 (2017), pp. 231–251 (cit. on p. 7).

[BH95] Yoav Benjamini and Yosef Hochberg. „Controlling the false discovery rate: a practical and powerful approach to multiple testing“. In: *Journal of the royal statistical society. Series B (Methodological)* (1995), pp. 289–300 (cit. on p. 17).

[Bir+09] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009 (cit. on pp. 22, 49).

[Bla00] Susan Blackmore. *The meme machine*. Vol. 25. Oxford Paperbacks, 2000 (cit. on pp. 6–8, 12, 13).

[Ble+03] David M Blei, Andrew Y Ng, and Michael I Jordan. „Latent dirichlet allocation“. In: *Journal of machine Learning research* 3.Jan (2003), pp. 993–1022 (cit. on p. 9).

[Bon87] Phillip Bonacich. „Power and centrality: A family of measures“. In: *American journal of sociology* 92.5 (1987), pp. 1170–1182 (cit. on p. 35).

[BY01] Yoav Benjamini and Daniel Yekutieli. „The control of the false discovery rate in multiple testing under dependency“. In: *Annals of statistics* (2001), pp. 1165–1188 (cit. on p. 17).

[Car+15] Samuel Carton, Souneil Park, Nicole Zeffer, et al. „Audience Analysis for Competing Memes in Social Media.“ In: *ICWSM*. 2015, pp. 41–50 (cit. on p. 6).

[CC13] David Chavalarias and Jean-Philippe Cointet. „Phylomemetic patterns in science evolution—the rise and fall of scientific fields“. In: *PloS one* 8.2 (2013), e54847 (cit. on pp. 6, 8, 33).

[CF07] Nicholas A Christakis and James H Fowler. „The spread of obesity in a large social network over 32 years“. In: *n engl j med* 2007.357 (2007), pp. 370–379 (cit. on pp. 6, 12, 13).

[CF08] Nicholas A Christakis and James H Fowler. „The collective dynamics of smoking in a large social network“. In: *New England journal of medicine* 358.21 (2008), pp. 2249–2258 (cit. on pp. 6, 12, 13).

[CF13] Nicholas A Christakis and James H Fowler. „Social contagion theory: examining dynamic social networks and human behavior“. In: *Statistics in medicine* 32.4 (2013), pp. 556–577 (cit. on pp. 6, 12, 13).

[CY06] Jingchun Chen and Bo Yuan. „Detecting functional modules in the yeast protein–protein interaction network“. In: *Bioinformatics* 22.18 (2006), pp. 2283–2290 (cit. on p. 9).

[Daw76] Richard Dawkins. *The selfish gene*. Oxford university press, 1976 (cit. on pp. 5, 7).

[Dee+90] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. „Indexing by latent semantic analysis“. In: *Journal of the American society for information science* 41.6 (1990), p. 391 (cit. on p. 9).

[Fre77] Linton C Freeman. „A set of measures of centrality based on betweenness“. In: *Sociometry* (1977), pp. 35–41 (cit. on p. 35).

[Hag+08] Aric Hagberg, Pieter Swart, and Daniel S Chult. *Exploring network structure, dynamics, and function using NetworkX*. Tech. rep. Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008 (cit. on p. 35).

[HC09] Francis Heylighen and Klaas Chielens. „Evolution of culture, memetics“. In: *Encyclopedia of complexity and systems science*. Springer, 2009, pp. 3205–3220 (cit. on p. 7).

[Jeo+01] Hawoong Jeong, Sean P Mason, A-L Barabási, and Zoltan N Oltvai. „Lethality and centrality in protein networks“. In: *Nature* 411.6833 (2001), p. 41 (cit. on p. 9).

[Jor+06] Ferenc Jordán, Wei-chung Liu, and Andrew J Davis. „Topological keystone species: measures of positional importance in food webs“. In: *Oikos* 112.3 (2006), pp. 535–546 (cit. on p. 9).

[Kat53] Leo Katz. „A new status index derived from sociometric analysis“. In: *Psychometrika* 18.1 (1953), pp. 39–43 (cit. on p. 35).

[KK16] MJ Krawczyk and K Kułakowski. „Authors as vehicles of scientific memes“. In: *International Journal of Modern Physics C* 27.09 (2016), p. 1650110 (cit. on p. 7).

[Kle99] Jon M Kleinberg. „Authoritative sources in a hyperlinked environment“. In: *Journal of the ACM (JACM)* 46.5 (1999), pp. 604–632 (cit. on p. 35).

[Kuh+14] Tobias Kuhn, Matjaž Perc, and Dirk Helbing. „Inheritance patterns in citation networks reveal scientific memes“. In: *Physical Review X* 4.4 (2014), p. 041036 (cit. on pp. 7, 12–15, 21).

[Les+09] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. „Meme-tracking and the dynamics of the news cycle". In: *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM. 2009, pp. 497–506 (cit. on p. 6).

[MS99] Christopher D Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. 1999 (cit. on p. 17).

[Pag+99] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. *The PageRank citation ranking: Bringing order to the web.* Tech. rep. Stanford InfoLab, 1999 (cit. on p. 35).

[Rom+11] Daniel M Romero, Brendan Meeder, and Jon Kleinberg. „Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter". In: *Proceedings of the 20th international conference on World wide web*. ACM. 2011, pp. 695–704 (cit. on p. 41).

[Sch02] Edna Schechtman. „Odds ratio, relative risk, absolute risk reduction, and the number needed to treat—which of these should we use?" In: *Value in Health* 5.5 (2002), pp. 431–436 (cit. on p. 25).

[Ser+09] M Ángeles Serrano, Marián Boguná, and Alessandro Vespignani. „Extracting the multiscale backbone of complex weighted networks". In: *Proceedings of the national academy of sciences* 106.16 (2009), pp. 6483–6488 (cit. on p. 38).

[Shi13] Limor Shifman. *Memes in digital culture*. 2013 (cit. on p. 6).

[SJ72] Karen Sparck Jones. „A statistical interpretation of term specificity and its application in retrieval". In: *Journal of documentation* 28.1 (1972), pp. 11–21 (cit. on p. 9).

[Sma73] Henry Small. „Co-citation in the scientific literature: A new measure of the relationship between two documents". In: *Journal of the Association for Information Science and Technology* 24.4 (1973), pp. 265–269 (cit. on p. 35).

[SO+07] Robert P St Onge, Ramamurthy Mani, Julia Oh, et al. „Systematic pathway analysis using high-resolution fitness profiling of combinatorial gene deletions". In: *Nature genetics* 39.2 (2007), p. 199 (cit. on p. 9).

[Sue+13] Caroline Suen, Sandy Huang, Chantat Eksombatchai, Rok Sosic, and Jure Leskovec. „Nifty: a system for large scale information flow tracking and clustering". In: *Proceedings of the 22nd international conference on World Wide Web*. ACM. 2013, pp. 1237–1248 (cit. on p. 6).

[TR12] Oren Tsur and Ari Rappoport. „What's in a hashtag?: content based prediction of the spread of ideas in microblogging communities". In: *Proceedings of the fifth ACM international conference on Web search and data mining*. ACM. 2012, pp. 643–652 (cit. on p. 6).

[TR15] Oren Tsur and Ari Rappoport. „Don't Let Me Be# Misunderstood: Linguistically Motivated Algorithm for Predicting the Popularity of Textual Memes." In: *ICWSM*. 2015, p. 426 (cit. on p. 6).

[Vie08] Anthony J Viera. „Odds ratios and risk ratios: what's the difference and why does it matter?" In: *Southern medical journal* 101.7 (2008), pp. 730–734 (cit. on p. 25).

[Wan+09] Yi Wang, Hongjie Bai, Matt Stanton, Wen-Yen Chen, and Edward Y Chang. „Plda: Parallel latent dirichlet allocation for large-scale applications“. In: *International Conference on Algorithmic Applications in Management*. Springer. 2009, pp. 301–314 (cit. on p. 9).

[Wen+14] Lilian Weng, Filippo Menczer, and Yong-Yeol Ahn. „Predicting Successful Memes Using Network and Community Structure.“ In: *ICWSM*. 2014 (cit. on p. 6).

[Özg+08] Arzucan Özgür, Thuy Vu, Güneş Erkan, and Dragomir R Radev. „Identifying gene-disease associations using centrality on a literature mined gene-interaction network“. In: *Bioinformatics* 24.13 (2008), pp. i277–i285 (cit. on p. 9).

List of Figures

List of Tables

Colophon

This thesis was typeset with $\LaTeX\,2_\varepsilon$. It uses the *Clean Thesis* style developed by Ricardo Langner. The design of the *Clean Thesis* style is inspired by user guide documents from Apple Inc.

Download the *Clean Thesis* style at `http://cleanthesis.der-ric.de/`.