

Models of human navigation in information networks based on decentralized search

Denis Helic
Knowledge Management
Institute
Graz University of Technology
Graz, Austria
dhelic@tugraz.at

Markus Strohmaier
Knowledge Management
Institute and Know-Center
Graz University of Technology
Graz, Austria
markus.strohmaier@tugraz.at

Michael Granitzer
Chair of Media Informatics
University of Passau
Passau, Germany
Michael.Granitzer@uni-
passau.de

Reinhold Scherer
Institute for Knowledge
Discovery
Graz University of Technology
Graz, Austria
reinhold.scherer@tugraz.at

ABSTRACT

Models of human navigation play an important role for understanding and facilitating user behavior in hypertext systems. In this paper, we conduct a series of principled experiments with *decentralized search* - an established model of human navigation in social networks - and study its applicability to information networks. We apply several variations of decentralized search to model human navigation in information networks and we evaluate the outcome in a series of experiments. In these experiments, we study the validity of decentralized search by comparing it with human navigational paths from an actual information network - Wikipedia. We find that (i) navigation in social networks appears to differ from human navigation in information networks in interesting ways and (ii) in order to apply decentralized search to information networks, stochastic adaptations are required. Our work illuminates a way towards using decentralized search as a *valid model* for human navigation in information networks in future work. Our results are relevant for scientists who are interested in modeling human behavior in information networks and for engineers who are interested in using models and simulations of human behavior to improve on structural or user interface aspects of hypertextual systems.

Categories and Subject Descriptors

H.5.4 [Information Interfaces and Presentation]: Hypertext/ Hypermedia—Navigation; H.5.3 [Information In-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

24th ACM Conference on Hypertext and Social Media
1–3 May 2013, Paris, France
Copyright 2013 ACM 978-1-4503-1967-6/13/05 ... \$ 15.00

terfaces and Presentation]: [Group and Organization Interfaces—Web-based interaction]

General Terms

Experimentation, Measurement, Algorithms

Keywords

Navigation, Decentralized search, Exploration, Exploitation

1. INTRODUCTION

In this work, we study the ways in which an established model for navigating social networks - i.e. *decentralized search* - can be applied to model human navigation in information networks.

Towards that end, we implement decentralized search as a model of navigation and evaluate it with a log of human navigational paths obtained from an online navigation game based on Wikipedia (TheWikiGame¹). In this game, a user starts at a randomly selected Wikipedia article (a so-called start page) and is supposed to navigate to another randomly selected article (a so-called target page) by following textual hyperlinks emanating from the current Wikipedia article only. During the game, each click is captured and stored in a log. We use a log of ca. 250,000 click paths to expose users' navigation behavior and to study factors that influence users' decision making during navigation.

In our experiments, we build on insights from Milgram's "small world" experiment [14] and on models of navigability; in particular decentralized search by J. Kleinberg [9], which we apply to model users' navigation. In addition, we apply models and approaches based on hidden metric spaces introduced by Boguna et al. [2], and explore how users' navigation behavior on Wikipedia can be modeled with decentralized search using topic hierarchies as background knowledge. The topic hierarchies may be constructed algorithmically.

¹<http://www.thewikigame.com>

mically from structural properties of the Wikipedia article network or external metadata about Wikipedia articles.

Research questions: In this paper, we are particularly interested in exploring answers to the following two main questions: (i) *Is decentralized search* - a model originally developed for social networks - *also applicable to information networks?* (ii) *And if it is, in what ways does it apply?* i.e. what variation of decentralized search best explains human navigation behavior on Wikipedia?

To answer these questions, we study different variants of decentralized search using different mechanisms for action selection. Action selection mechanisms for decentralized search can be grouped into two main categories, i.e. greedy (deterministic) and stochastic greedy (probabilistic) approaches. Within the stochastic greedy mechanism, we will investigate three probabilistic action selection strategies, i.e. ϵ -greedy, softmax rule, and inverse distance rule. Using these models of navigation, we generate synthetic click paths and compare them to actual click paths generated by users.

Contributions: Our work makes the following contributions: (i) We present and discuss several potential differences and commonalities between navigation in social and information networks. (ii) We identify certain configurations of decentralized search that are capable of modeling human navigation behavior in information networks in valid ways. (iii) We introduce a *novel action selection mechanism* for decentralized search in information networks called *decaying ϵ -greedy*. This new mechanism is general enough to be adapted to different kinds of navigation scenarios and best explains the observed parameters of navigation in our dataset.

More specifically, we find that human navigation in information networks is a two phase process combined of the *exploitation of the known* and the *exploration of the unknown*, as opposed to search in social networks which largely exploits the local knowledge of decentralized agents. Our results suggest that humans either follow specific links on purpose (exploitation) whenever they are confident enough that those links bring them closer to their particular target, or they select links almost arbitrarily at random (exploration) whenever they do not possess enough knowledge to relate the candidate links to their target. In particular, we empirically measure the ratio between exploitation and exploration in human navigational paths and we find that in our dataset, *exploration accounts for up to 15-20% of clicks and exploitation accounts for 80-85%* on average. These results are strikingly similar to the α parameter used in Brin and Page's PageRank where α is set to 15%. We also investigate the relation of this ratio to users' progress during navigation and find that the ratio of exploration decays as navigation progresses. In other words, users tend to explore more in the early phases of navigation games and tend to exploit more in the later navigation phases as they get more confident about their position in the network.

The paper is structured as follows: In section 2 we review the state-of-the-art in this domain. Next we present our methodology and experimental setup in section 3 and discuss our results in sections 4, 5, and 6. Section 7 concludes our paper.

2. NAVIGATION IN NETWORKS

Research on (decentralized) search in social networks was

initiated by the small-world experiment conducted by Milgram [14]. In that experiment, randomly selected persons from several locations in the US were required to pass a letter to a target person in Boston in a decentralized manner, i.e. through their social networks. One of the most interesting results of that experiment was that the average length of the letter chain was only six, and thus the entire social network in the US represented a *small-world*. Apart from the findings that people in a social network are connected by short chains of acquaintances, another result of this experiment received a lot of attention from various researchers. Not only are humans connected by short chains, but they are also able to efficiently find these short chains in a *decentralized* manner. For example, Kleinberg concluded that humans possess *background knowledge* of the network structure and that this knowledge allows humans to efficiently find short paths [8, 9, 10]. Kleinberg represented such background knowledge as a hierarchy of nodes, where more similar nodes are situated closer to each other in the hierarchy. Independently, Watts [20] introduced the notion of social identity as a membership in a number of social groups organized in hierarchies and showed the existence of efficient decentralized search algorithms that utilize these hierarchies by simulation. In [1], Adamic reports of the results of empirical studies investigating decentralized search in social networks. Adamic simulated search in real-world networks such as an organizational e-mail network and an online student network. The simulations were based on Kleinberg's decentralized search algorithm with different hierarchies being applied as background knowledge, e.g., an organizational hierarchy as well as a hierarchy reflecting the position of a person in the physical space. Results showed that both of these hierarchies can be effectively used to support decentralized search. In addition, Adamic has shown that the performance of a decentralized search algorithm depends on the quality of the hierarchical background knowledge and therefore the hierarchies that are used for guiding the decentralized search play an important role in those network navigation models. In summary, decentralized search represents a very well studied theory that can explain navigation in social networks since the theory has been validated in several experimental studies after its introduction.

Crucial to all variants of decentralized search is the notion of *distance* between network nodes, i.e. there exists a metric defined on a set of network nodes. According to this distance function, the decentralized search algorithms are *greedy*, i.e. an agent always selects a node with the smallest distance to a particular target node. Accordingly, those algorithms are also *deterministic* because the algorithm always produces the same navigation sequence between a given start and a given target node.

In [18], the authors further investigate the notion of the node distance function by discussing *hidden metric spaces*. In such hidden metric spaces, nodes are identified by their co-ordinates - distance between nodes is their geometric distance in a particular metric space. To that effect, hierarchical background knowledge as originally introduced by Kleinberg may be comprehended as a metric space where the distance is defined as e.g. the shortest path between nodes in the hierarchy.

One of the original Kleinberg's suggestions was that decentralized search also represents an intuitive model of navigation in information networks, e.g. on the Web. Similarly

	<i>Social Networks</i>	<i>Information Networks</i>
Agents per search	multiple agents	single agent
Type of routing	decentralized (with local knowledge)	centralized (with local knowledge)
Searcher	part of the network (endogenous)	not part of the network (exogenous)
Routing decisions	social intuitions	topical intuitions
Local knowledge	rich	limited
Consultation of candidates	costly	cheap

Table 1: Potential differences and commonalities between navigation in social and information networks

to search in social networks, navigation in information networks is inherently *local*, i.e. when users navigate they are only aware of the links emanating from the document currently displayed at their screens. One can assume that they apply similar greedy strategies in their navigation towards the destination document. For example, users might tap into their intuitions about *conceptual* similarities and distances between documents and select the document that has conceptually the smallest distance to the destination document. Based on these ideas, we have applied decentralized search as an evaluation strategy for estimating the quality of a hierarchy for supporting user navigation in our own previous work. For example, in [6] we evaluated folksonomy construction algorithms from a pragmatic perspective, i.e. we evaluated their efficiency for supporting navigation in social tagging systems. More recently, in [19] we made first steps towards validation of decentralized search as a model of navigation in information networks by comparing deterministic greedy decentralized search with human click paths in a Wikipedia navigation game. One of the conclusions was that decentralized search with appropriately constructed hierarchical background knowledge is more efficient than humans when they navigate, as decentralized search finds shorter paths to randomly selected destination nodes on the average as compared to humans.

In other related work, West and Leskovec [22] performed a study of user navigation behavior. The authors analyzed a collection of click paths of users who were playing a navigation game in a network of links between the concepts of Wikipedia. In their work, they found that user navigation behavior differs from shortest paths. For example, users typically navigate through high-degree hubs in the early phase and then apply content similarity as a criteria for finding the destination node. In subsequent work [21], the authors analyzed a number of decentralized search algorithms using various distance functions and benchmarked them against their human click corpus. The authors also found that even simple search strategies such as utilizing node degrees outperform human information seeking.

Although these initial findings are very promising and show that decentralized search may indeed be a first step in modeling navigation in information networks, they have also revealed subtle differences between models of navigation in social and information networks. For example, simple deterministic greedy model of navigation in a typical case outperforms human navigation in terms of navigation efficiency. We take this initial observations as a starting point for studying the differences between navigation in social and information networks in greater detail.

Table 1 highlights selected differences between navigation in social and information networks.

Navigation in social networks (c.f. the Milgram ex-

periment [14]) usually involves *multiple agents* engaging in a collective search effort where every agent tries to forward a request (e.g. a letter) to other agents who are presumably closer to a given target person. In this sense, navigation in social networks represents *decentralized routing*, where multiple agents collectively make routing decisions in sequential order. The agents themselves are also nodes in the social network which they are navigating (in other words: the agents are *endogenous* to the network). Routing decisions are based on *social intuitions* about which other agent most likely moves the search closer to the target node. At each step, the corresponding agent has access to *local knowledge* about the network *only*, i.e. knowledge about the nodes that represent her ego-network (the one-hop neighborhood). Typically, each agent has *rich knowledge* about her local social neighborhood, and she can tap into her intuitions to contact the right candidate node to move the search forward. However, consulting candidate nodes [12] is usually *costly* - one can easily assert that exhaustively contacting all candidates from the set of candidates would not represent a feasible choice.

Navigation in information networks (e.g. navigation on Wikipedia) usually involves just a single agent who tries to navigate to target pages. In this sense, navigation in information networks represents *centralized routing*, where navigation decisions are made by an autonomous agent. The agent herself is not part of the information network, she merely navigates it (in other words: the agents are *exogenous* to the network). Routing decisions are based on *topical intuitions* about which page most likely moves the searcher closer to the target page. At each step, the agent only has *local knowledge* about the network, i.e. knowledge about the links that are emanating from a particular page (the one-hop neighborhood). Typically, an agent who navigates an information network has *limited knowledge* about the local neighborhood for a given node. Yet, the costs for consulting candidate nodes is *low*, and we can easily see that exploration of candidate nodes and backtracking to existing nodes would represent a feasible navigation strategy.

Summary: While navigation in social and information networks share some interesting properties (e.g. local knowledge of the network, navigation based on intuitions), a number of interesting differences can be observed. Overall, these differences suggest that while greedy approaches to modeling navigation in social networks appear reasonable, in information networks they might need to be adapted because of the stochastic nature of human navigation behavior [7]. These observations make it interesting to study the ways in which models developed for navigation in social networks might be applied to modeling navigation in information networks. We will outline our methodology towards answering this question and our experimental setup next.

3. METHODOLOGY

We adopt decentralized search as a model of navigating social networks and apply it to Wikipedia. While the original proposal for decentralized search is based on a greedy approach, we will investigate the utility of different action selection mechanisms. We know from previous research that navigation on information networks such as the World Wide Web exhibits high variation [7] which can not be captured by a deterministic greedy modeling approach. We also know that such navigation can be modeled via stochastic processes. By experimenting with different action selection mechanisms, we want to inject different kinds of stochasticity and find out which one explains actual user behavior best.

3.1 Action selection mechanisms

We will experiment with the following action selection mechanisms presented in Table 2:

Greedy (baseline): With greedy action selection, the algorithm always chooses the candidate node j with the minimal distance $d(j, t)$ to the target node.

ϵ -greedy: With ϵ -greedy action selection, the algorithm chooses the candidate node j with the minimal distance $d(j, t)$ to the target node with $1-\epsilon$ probability, and with probability ϵ it chooses another candidate uniformly at random.

Softmax rule: With softmax [3, 5], the algorithm chooses a candidate node j with probability $p(j) \propto e^{cf(j)}$, where $f(j)$ represents a fitness function calculated from the distances $d(j, t)$, and c corresponds to an agent's confidence in her intuitions. When c is large enough (high confidence), softmax always chooses the candidate node j with the minimal distance $d(j, t)$, which makes it equivalent to greedy action selection. However, with small values of c (low confidence), the algorithm tends to select other candidate nodes based on $f(j)$.

Inverse distance rule: With inverse distance action selection (c.f. [13]), the algorithm chooses a candidate node j with $p(j) \propto f(j)^{-c}$, where c again represents a confidence parameter. The properties of this rule are similar to the softmax rule - the only difference lies in the probability distribution for action selection.

3.2 Experimental Setup

TheWikiGame dataset contains almost 800,000 navigation games. Of these, more than 250,000 are successful games, i.e. games where users have reached the target page successfully. In this paper, we concentrate on the analysis of the successful games only. To obtain the underlying information network, we use the April 2012 Wikipedia dump. In this network, Wikipedia articles are represented as nodes and hyperlinks between Wikipedia articles are represented as links. The network contains around 10,000,000 nodes and 250,000,000 edges.

For simulation, we use decentralized search with hierarchical background knowledge where action selection mechanisms are configured according to Table 2. As hierarchical background knowledge, we construct hierarchies from the Wikipedia network using the algorithm introduced in [15]. The algorithm is based on the idea that many networks tend to be inherently hierarchically structured. Such hierarchical structure typically leads to the emergence of observable structural properties such as power-law degree distributions and high node clustering (cf. [4]). The authors develop their ideas by suggesting that each link in a network is either of

a hierarchical type or some other type, e.g. it is a synonym, or a general association. The algorithm aims to recognize these hierarchical links - it iterates through all links in the network and decides, using a simple criteria, if that link is of a hierarchical or some other type. The hierarchical links are kept in the network and all other links are removed from the network. The algorithm decisions are based on a so-called hierarchical score, which is a measure of the generality of a node. For each link a hierarchical ratio between hierarchical scores of two incident nodes is calculated. If the hierarchical ratio is close to 1 then those two nodes are close in generality and they are situated in the same hierarchy level - thus, the link between those two nodes is not a hierarchical one and is therefore removed from the network. Similarly, if the hierarchical ratio for a link is close to 0, then those two nodes are very far away from each other in the hierarchy and the link is also removed. Technically, the authors define two thresholds - a high and a low threshold - to decide on link removal. Thus, a link is removed if the hierarchical ratio is greater than the high threshold or smaller than the low threshold. In our experiments, we use a local flow score as the hierarchical score, which is defined as:

$$g(i) = \frac{k_i^-}{k_i^+} \sqrt{k_i^-}, \quad (1)$$

where k_i^- is the in-degree of node i , and k_i^+ is the out-degree of node i . The term $\sqrt{k_i^-}$ ensures that a node having e.g. 2000 in-degree and 1000 out-degree is rendered more general than a node having e.g. 2 in-degree and 1 out-degree. Next, for each link (i, j) we calculate the hierarchical ratio as:

$$r(i, j) = \frac{\min(g(i), g(j))}{\max(g(i), g(j))}. \quad (2)$$

If the $r(i, j)$ falls within two thresholds then the link (i, j) is kept in the network, otherwise it is removed. As thresholds we choose 0.9 and 0.1 for high and low thresholds respectively (cf. [15]) and extract the hierarchy from the Wikipedia network. The hierarchy contains around 5,000,000 nodes and 60,000,000 links.

Next, we apply the constructed hierarchy as the background knowledge and initialize these different versions of decentralized search with pairs of start and target nodes. The number of simulations per configuration corresponds to the number of successful games (more than 250,000 games) played on TheWikiGame.

For evaluation, we calculate a series of measures to characterize the ability of different action selection mechanisms to mimic human behavior. First, we calculate *success rates*:

$$s = \frac{|W|}{|P|}, \quad (3)$$

where W is the set of node pairs which the simulator successfully played and P is the set of all node pairs used for simulation. Thus, the success rate captures the percentage of cases where humans respectively decentralized search were successful in finding the target. A success rate of 0.9 means that in 90% of cases, the searcher did find the target. In addition to the success rate we calculate *stretch*, which is a measure for the efficiency of navigation. It is calculated by dividing the lengths of the actual paths $h(s, t)$ by the lengths of the shortest paths $l(s, t)$ between start and target

Action selection	Definition	Description
Greedy	$j = \min(M \setminus C)$	An agent who always follows her intuitions
ϵ -greedy	$j = \min(M \setminus C)$, with $P = 1 - \epsilon$ $p(j) = \frac{1}{ \Gamma(i) }$, with $P = \epsilon$	An agent who follows her intuitions with probability $1 - \epsilon$, but who adopts a random strategy with probability ϵ
Softmax rule	$p(j) = \frac{e^{cf(j)}}{\sum_j e^{cf(j)}}$, $f(j) = 1 - \frac{d(j,t)}{\max_{k,l \in V} d(k,l)}$	An agent who follows her intuitions to different extents. The parameter c sets the extent to which the agent is greedy or stochastic according to her intuitions.
Inverse distance rule	$p(j) = \frac{f(j)^{-c}}{\sum_j f(j)^{-c}}$, $f(j) = \frac{d(j,t)}{\max_{k,l \in V} d(k,l)}$	An agent who follows her intuitions to different extents. The parameter c sets the extent to which the agent is greedy or stochastic according to her intuitions.

Table 2: Different action selection mechanisms for Decentralized Search. The following definitions apply: V is the set of all nodes that is totally ordered under the binary relation $<$ (e.g. node encodings may be sorted), t is the target node, i is the current node, $\Gamma(i)$ is the set of all neighbors of i , $j \in \Gamma(i)$ is a candidate node at node i , $d(j,t)$ is a distance between nodes j and t . Further, $M := \operatorname{argmin}_{j \in \Gamma(i)} d(j,t)$ is the argument of the minimum of the distance function of candidate nodes at node i , and finally C is the set of already visited nodes.

nodes and then averaging over all nodes:

$$\tau = \frac{1}{|W|} \sum_{s,t \in W} \frac{h(s,t)}{l(s,t)}. \quad (4)$$

A stretch $\tau = 2$ means that on average, actual paths lengths are twice as large as the corresponding shortest paths.

Next, we analyze the distribution of the hop lengths of different simulation variants and compare those to the shortest path distribution and human hop length distribution. We quantify the differences between distributions by calculating Kullback-Leibler divergence D_{KL} [11]. This measure tells us how different the distributions generated by the models are from the distributions of actual click path. In information theory D_{KL} measures the number of extra bits that is needed to code the real distribution with the codes derived from the model distribution. Thus, lower values of D_{KL} represent better approximations of the real distribution, i.e. the model models the reality more closely. Kullback-Leibler divergence is calculated as:

$$D_{KL}(p||q) = \sum_x \log_2 \left(\frac{p(x)}{q(x)} \right) p(x) \quad (5)$$

Typically, we say $D_{KL}(p||q)$ is the Kullback-Leibler divergence of q from p , where p is the real distribution and q is the model distribution.

4. RESULTS

We are organizing the presentation of our results around the two leading research questions of this paper:

(i) *Is decentralized search* - originally developed in the context of social networks - *also applicable to information networks?*

Figure 1, 2 and 3 show the results from our experiments. Figure 1 depicts success rates and stretch depending on different action selection mechanisms. While success rates seem to be mostly insensitive to the different mechanisms

and parameters, we find that stretch is highly sensitive to the choice of the parameter. For example, from Figure 1a we can see that a greedy decentralized search model (with $\epsilon = 0$) does *not* explain human click paths very well. With an $\epsilon = 0.15$, the stretch generated by ϵ -greedy matches the stretch of human click paths much better (depicted by the intersection of the green and the pink line). These results suggest that greedy search - which was shown to explain navigation in social networks well - does not fully explain navigation in information networks. In particular, navigation in information networks seems to be a more stochastic process where phases of deterministic greedy behavior are seamlessly combined with random probabilistic behavior. Thus, during deterministic phases, users seem to have a high confidence into their intuition about the network and *exploit* them in a greedy manner. This helps to reduce the conceptual distance to the target page at each step. On the other hand, during random phases, users do not possess sufficient intuitions or are not confident enough in their intuitions about the network and therefore switch to *exploration* of their immediate neighborhood, apparently selecting links on random until they arrive at a page where their intuitions can again lead them towards the target page, and they switch again to the exploitation phase. Concerning the applicability of decentralized search, we can conclude that injecting certain levels of randomness in the baseline (greedy) decentralized search models seems to produce global navigation properties such as success rate and stretch that are close to human performance.

(ii) *In what ways does decentralizes search apply?* i.e. what variation of decentralized search best explains human navigation behavior on Wikipedia?

To answer the second research question we turn to a more detailed analysis of the properties of simulated navigation. We are now interested in the underlying distributions that generate the global properties discussed previously. Thus, Figure 2 depicts the hop length distributions of different ac-




Figure 1: Success rate and stretch of different navigation models. Across all action selection mechanisms (ϵ -greedy, Softmax, Inverse distance rule) we can see that the success rate of humans and decentralized search is comparable (close to 1.0, see the two lowest horizontal lines above). Interesting differences emerge when comparing the stretch of human clicks (upper horizontal line) with the stretch of different action selection mechanisms. Subfigure 1a shows that with an $\epsilon = 0.15$, human stretch can be approximated well. Greedier strategies (lower epsilons) tend to reduce stretch to the extent that decentralized search significantly outperforms humans, while more random strategies (higher epsilons) tend to increase stretch. The same holds - with different parameters and to different extents - for Softmax 1b and the Inverse distance rule 1c.




Figure 2: Hop length distributions generated by different navigation models. Across all action selection mechanisms (ϵ -greedy, Softmax, Inverse distance rule) we can observe that our models tend to produce hop lengths distributions (dashed lines) that coarsely approximate the hop length distribution of human paths (solid green line). The hop length distribution of the shortest paths (solid red line) shows the topological baseline. While the distributions of models and human paths are similar, we can also observe that the models tend to be more efficient, i.e. on average, they find the targets faster than humans. For example in 2a, ϵ -greedy with $\epsilon = 0.1$ or 0.15 exhibits a hop length distribution very close to the hop length distribution of humans, while $\epsilon = 0.5$ or $\epsilon = 0$ (i.e. greedy decentralized search) outperform humans. Although some of the hop length distributions of the simulator (c.f. the quantitative analysis with Kullback-Leibler divergence in Figure 3) are very close to the hop length distribution of the humans one specific difference can be observed. The human hop length distribution is shifted to the right by 1, i.e. the human distribution has its mode at 4, whereas all simulator distributions have their modes at 3. Similar trends can be observed across different action selection mechanisms 2b and 2c.

tion selection mechanisms. Their differences are quantified by the Kullback-Leibler divergence between human distributions and model distributions as depicted in Figure 3. A striking result from this analysis is that in a typical case, the divergence between simulator hop length distribution and human hop length distribution is only minimal. For example, ϵ -greedy strategy with an $\epsilon = 0.15$ has $D_{KL} = 0.22$. We can interpret this result in the following information-theoretic way: if we would encode human hop lengths with the code defined by the ϵ -greedy strategy with the $\epsilon = 0.15$, we would on average need only 0.22 bits more than if we

would encode the human hop lengths with the code defined by the real distribution.

Figure 3 shows what combination of action selection mechanism and parameter best approximates the hop lengths distributions of human click paths. We see that a combination of ϵ -greedy with an $\epsilon = 0.15$ minimizes the Kullback-Leibler divergence of model distributions from human distributions. This means that ϵ -greedy best explains human navigation with regard to hop lengths behavior. Intuitively, we can say that a user tends to exploit her intuitions (background




Figure 3: Similarity of hop length distributions between models & humans, and models & shortest path. Using the hop length distributions from Figure 2, we calculate the Kullback-Leibler divergence of model distributions from human distributions (green line) and model distributions from shortest path distributions (red line) depending on different action selection mechanisms and parameters. This explains what parameters produce hop length distributions that are closest to the hop length distributions of humans. In 3a we can see that for $\epsilon = 0.1, 0.15$ or 0.2 , the KL divergence is minimal, which means that the models represent a good fit with human click behavior. A similar trend - with different parameters - can be observed for Softmax 3b and the Inverse distance rule 3c.

knowledge) 85% of the time, and explores new options 15% of the time.

Another interesting result is that the simple ϵ -greedy strategy performs even slightly better than more complex action selection strategies such as softmax or inverse distance rule. We believe that the reason for this is that human navigation behavior can be fairly well modeled as a discrete process combined of exploration and exploitation phases. In other words, humans either explore or exploit the network, whereas softmax and inverse distance rule suggest that humans always possess certain intuitions about the network but with varying confidence levels at different stages of navigation.

5. A NOVEL MECHANISM FOR ACTION SELECTION

Although these first experimental results are very promising, we can also observe an interesting deviation of the simulator hop length distributions from human hop length distribution. While the mode of the human hop length distribution is at 4, all simulator distributions have their mode at 3. This is likely caused by the fact that humans rarely find the target node as early as the simulator (e.g in 2 steps). In fact, humans are able to find the target node only in 5.7% of cases in two hops in our dataset, whereas the best performing action selection strategy ϵ -greedy with $\epsilon = 0.15$ does the same in 17.3% cases. Thus, the simulator is by large more efficient in early navigation phases. This leads us to the conclusion that humans do not possess sufficient intuitions in the beginning of the navigation game, and as the game progresses, their intuitions become better which allows them to close in on the target node. We believe that there may be two reasons for such a behavior.

First, as the authors of [22] found that the early phases of human navigation are dominated by node degrees – humans tend to reach a high-degree node very quickly and then continue from there by utilizing the node similarity until they reach the target. A possible explanation for the human inefficiency in the first few navigation steps may be that the




Figure 4: Rates of exploration over degrees (left) and hop length (right). The relation between exploration probability and node degree seems to be non-linear, i.e. exploration probability is not proportional to node degree as our experiments confirm (Figure 4a). On the other hand, Figure 4b clearly shows that exploration rate drops with each new click as humans orient themselves better in the network.

willingness to explore is *proportional to the node degree*, e.g. humans possess fewer intuitions and less confidence in their intuitions on pages with more links and therefore, they explore more on such pages.

Second, humans may need a few initial clicks to orient themselves in the network and therefore explore more in early phases. This is an intuitive explanation because the navigation game setup lets humans start on a *randomly* selected page, which is typically not related to the target page. Thus, at the beginning of navigation, humans tend to *explore* the network to a larger extent, and once when they have oriented themselves in the network, they tend to be more exploitative in their search for the target node.

To test those two hypothesis we investigate how exploration and exploitation depend on the node degree and the current hop length, i.e. how those two features help in classification of a click into an exploitation or an exploration class. To that end, let i be the current node in a naviga-

tion sequence. We define a click on node j to belong to the exploitation class if node j has the minimal distance to the target node, i.e. $j \in M$, where $M := \operatorname{argmin}_{j \in \Gamma(i)} d(j, t)$. Otherwise if the distance of the selected node j to the target node is greater than the minimal distance we define that click as exploration click.

Figure 4 shows the exploration probabilities of a click dependent on node degrees and the current hop length. The relation between node degree and exploration probability seems to be a non-linear one – a simple model of exploration that is proportional to the node degree can not capture that relation. Indeed, we performed a few experiments where we set varied ϵ in proportion with node degree and could not observe any improvement in our measures in comparison with the standard ϵ -greedy action selection. On the other hand, relation between hop length and exploration probability is much simpler and follows our previous intuition. Therefore we turn our attention to modeling ϵ as a function of the number of pages that users visited so far.

Decaying ϵ -greedy action selection: To further improve the validity of the models based on these observations, we propose a new action selection mechanism which we call *decaying ϵ -greedy* model. The mechanism is based on a decay function that adapts ϵ at every step during navigation. The decay function can take numerous forms, in this work we experimented with a decay function that starts with a given ϵ_0 (e.g. $\epsilon_0 = 0.8$) and then decays during navigation by a certain factor λ (e.g. $\lambda = 2$). In general, we define ϵ as the function of the hop length t in the following way:

$$\epsilon(t) = \epsilon_0 \lambda^{-t}. \quad (6)$$

With the exemplary parameters from above, *decaying ϵ -greedy* would use the following epsilons ϵ_t at step t during navigation: $\epsilon(0) = 0.8$, $\epsilon(1) = 0.4$, $\epsilon(2) = 0.2$, and so on and so forth modeling the increase in human intuition about the network from click to click. Note that this approach is fairly general and can be easily adopted to other navigation settings. For example, suppose that a user arrives on Wikipedia from a search engine and lands on a page that is not completely random, i.e. it is - to a certain extent - already related to a given target page. Then, we can easily model this situation by selecting a lower value for ϵ_0 (e.g. 0.4) and account for the fact that the landing and target page are related.

As we see in Figure 5, the decaying ϵ -greedy strategy with $\lambda = 2$ further improves the validity of decentralized search, e.g. for $\epsilon_0 = 0.9$ we observe a perfect match of the success rate and stretch to those of the humans and with $D_{KL} = 0.05$ the distributions are almost identical.

For the next test that we perform we create a new hierarchy for the background knowledge. We apply the same algorithm as before but set new threshold, i.e. we set 0.6 and 0.3 for the high and low threshold, respectively. The newly constructed hierarchy posses fewer nodes (about 4,000,000) and links (about 20,000,000) as previously (because we apply a more strict criteria for hierarchical links and more links are therefore removed from the network). We observe similar behavior as before – the simulator produces almost a perfect match in success rate and stretch for $\epsilon_0 = 0.7$ and $\epsilon_0 = 0.8$ with $D_{KL} = 0.04$ and $D_{KL} = 0.03$, respectively (see Figure 6).

We have the following explanation for slightly lower values for ϵ_0 in this experiment – the hierarchy is smaller, and it is

more difficult for the simulator to find its way towards the target node if it drifts too much away from the greedy path in the exploration phase. When the hierarchy is larger, the simulator has a larger background knowledge which allows it to find its way back to an optimal path even from very far and distant places. Thus, this leads to the conclusion that, although ϵ_0 also depends on the size of the background knowledge the proposed mechanism is robust, general, and can account (with a proper parameter configuration) not only for different starting conditions but also for differences in the applied background knowledge.

6. DISCUSSION

Our results suggest that while simple (greedy) decentralized search represents an intuitive model for navigation in information networks, it does not fit human behavior well. In this section, we want to leverage the results from our experiments to revisit the differences between navigation in social and information networks laid out in Table 1. What are those differences exactly and how can they be explained?

Based on our experiments, we can draw the following conclusions:

Complete vs. Incomplete knowledge: In social networks, humans in general have good knowledge about their ego-network (the set of candidates that can directly be reached by them). They can exploit this knowledge when deciding who the best candidate for forwarding a letter would be. In information networks, humans might have limited or no knowledge about the links that are emanating from a particular document, and therefore their selection in many cases might be completely arbitrary. In our models we accounted for this situation by injecting randomness into the navigation process, i.e. with a probability $\epsilon(t)$ the models select a page from the candidate pages uniformly at random. Moreover, with ϵ decaying with each navigation step we account for humans increasing their orientation sense as they progress through the network. Also as our experiments suggest different starting pages and their relatedness to the target page, as well as the size of the background knowledge seem to play an important role in the navigation and the interplay between exploitation and exploration.

Exploration vs. Exploitation: We come to the conclusion that in information networks navigation is a combination of modes which consist of an orientation (or *exploration*) mode and an goal-seeking (or *exploitation*) mode. In other words, whenever the users lack the knowledge to identify the best candidate node they inevitably - because of low cost involved in the process - explore the network to find another node where their knowledge allows them to switch again to a goal-oriented exploitation mode of navigation. Thus, our experiments not only provide empirical evidence that human navigation is a combination of exploitation and exploration modes, they also provide an estimate of the ratio between these two modes. In our data, exploitation dominates exploration by a factor of ca. 1 : 7. The former is adopted whenever humans posses (or believe to posses) enough knowledge to select an optimal link that brings them closer to the destination, and the latter is adopted whenever humans are not sure about what link would bring them closer to the destination.

Similar observations have been made in past reserach, although on a more theoretical level. For example, in the framework of information foraging theory [17] Pirolli pos-




Figure 5: For a particular parameter configuration ($\lambda = 2$, $\epsilon_0 = 0.9$) success rate and stretch of the decaying ϵ -greedy match those of humans. The simulator exhibits more variance in the early navigation phases and is less efficient in the first few navigation steps than e.g. the standard ϵ -greedy. This renders the hop length distribution of the decaying ϵ -greedy more similar to the human hop length distribution, e.g. the hop length distribution for the best performing configuration is shifted to the right in a similar way as the human hop length distribution and has its peak at 4. This is also quantifiable by Kullback-Leibler divergence, which for the best configuration is only $D_{KL} = 0.05$.




Figure 6: With a smaller hierarchy decaying ϵ -greedy still matches humans to a great extent – the action selection mechanism is robust and can model different starting conditions, as well as differences in the background knowledge.

tulates that when users explore information structures they constantly assess the quality of available links and then follow the most promising links to satisfy their information need. The cues that indicate the possibility to find more and useful information are sometimes called *information scent* [16]. According to the theory user behavior in information systems is guided by the constant estimation of the cost and value of information structured in patches in respect to the current user information need. Principally, when the value of information is high then the users exploit information contained in the current patch and whenever that value is low and to the same time the cost of leaving the current patch is also low the users explore the available patches to find one with high information value.

Summary: In information networks, humans appear to act greedy if they think they have good intuitions about the next click (high confidence), and they act randomly if they think they lack good intuitions. The extent to which they lack good intuitions can be modeled with the parameters introduced (e.g. ϵ). In this work, our main interest was to validate the models with regard to success rate, stretch, and hop length distribution. We leave the validation of other

aspects of human navigation (e.g. comparing the extent to which the paths generated by models are semantically similar to the paths generated by humans) to future work.

7. CONCLUSIONS

To the best of our knowledge, this is the first work to investigate the extent to which decentralized search - a model that has been originally developed for navigation in social networks - can be applied to information networks.

The main contributions of this work are three-fold:

First, we identify several potential differences and commonalities between navigation in social and information networks. These differences suggest that in order to apply decentralized search to model human navigation in information networks, adaptations are necessary. This has led us to the design of a series of principled experiments that investigated the validity of different variations of decentralized search.

Second, we have conducted experiments geared towards comparing different action selection mechanisms and corresponding parameters for decentralized search. In this experiments, we have identified certain configurations of decentral-

ized search that are capable of modeling human navigation behavior in information networks in valid ways.

Third, we introduce a *novel action selection mechanism* for decentralized search in information networks called *decaying ϵ -greedy*. This model is based on a decay function that adapts ϵ at every step during navigation. We find that this new model is (i) general enough to be adapted to different kinds of navigation scenarios and (ii) best explains the observed parameters of navigation in our dataset.

Subsequent research can build on these contributions and use variations of decentralized search - and corresponding parameters that need to be adapted to given settings - as a valid model for human navigation. Thereby, our results are relevant for scientists who are interested in modeling human behavior in information networks and for engineers who are interested in using models and simulations of human behavior to improve on structural or user interface aspects of hypertextual systems.

In future work, we believe that it would be interesting to conduct comparative studies of navigation in a series of other information networks. This would allow us to - for example - understand the extent to which the identified parameters (e.g. ϵ_0 or λ) are specific to our Wikipedia setting, or are universal. Other future work could focus on quantifying the potential differences and commonalities between social and information networks in comparative empirical studies. Finally, we hope that our work sparks an interest to further develop and use decentralized search as a tool for modeling human behavior in hypertextual systems.

8. ACKNOWLEDGMENTS

This research was in part funded by the FWF Austrian Science Fund research project "Navigability of Decentralized Information Networks" (P24866).

9. REFERENCES

- [1] L. Adamic and E. Adar. How to search a social network. *Social Networks*, 27(3):187 – 203, 2005.
- [2] M. Boguñá, D. Krioukov, and K. C. Claffy. Navigability of complex networks. *Nature Physics*, 5:74–80, Jan. 2009.
- [3] J. S. Bridle. Probabilistic Interpretation of Feedforward Classification Network Outputs, with Relationships to Statistical Pattern Recognition. In Fogelman-Soulie and Herault, editors, *Neurocomputing: Algorithms, Architectures and Applications*, NATO ASI Series, pages 227–236. Springer, 1990.
- [4] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453(7191):98–101, 2008.
- [5] N. Daw, J. O’Doherty, P. Dayan, B. Seymour, and R. Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–879, 2006.
- [6] D. Helic, M. Strohmaier, C. Trattner, M. Muhr, and K. Lerman. Pragmatic evaluation of folksonomies. In *Proceedings of the 20th international conference on World wide web*, WWW ’11, pages 417–426, New York, NY, USA, 2011. ACM.
- [7] B. A. Huberman, P. L. T. Pirolli, J. E. Pitkow, and R. M. Lukose. Strong regularities in world wide web surfing. *Science*, 280(5360):95–97, 1998.
- [8] J. Kleinberg. The small-world phenomenon: an algorithm perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, STOC ’00, pages 163–170, New York, NY, USA, 2000. ACM.
- [9] J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845, August 2000.
- [10] J. M. Kleinberg. Small-world phenomena and the dynamics of information. In *Advances in Neural Information Processing Systems (NIPS) 14*, page 2001, Cambridge, MA, USA, 2001. MIT Press.
- [11] S. Kullback and R. A. Leibler. On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86, 1951.
- [12] E. Lebhar and N. Schabanel. Almost optimal decentralized routing in long-range contact networks. *Automata, Languages and Programming*, pages 73–96, 2004.
- [13] G. Miao, S. Tao, W. Cheng, R. Moulic, L. Moser, D. Lo, and X. Yan. Understanding task-driven information flow in collaborative networks. In *Proceedings of the 21st international conference on World Wide Web*, pages 849–858. ACM, 2012.
- [14] S. Milgram. The small world problem. *Psychology Today*, 1:60–67, 1967.
- [15] L. Muchnik, R. Itzhack, S. Solomon, and Y. Louzoun. Self-emergence of knowledge trees: Extraction of the wikipedia hierarchies. *Phys. Rev. E*, 76:016106, Jul 2007.
- [16] P. Pirolli. Computational models of information scent-following in a very large browsable text collection. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI ’97, pages 3–10, New York, NY, USA, 1997. ACM.
- [17] P. Pirolli and S. Card. Information foraging. *Psychological Review*, 106(4):643–675, 1999.
- [18] M. A. Serrano, D. Krioukov, and M. Boguñá. Self-similarity of complex networks and hidden metric spaces. *Phys. Rev. Lett.*, 100(7):078701, Feb 2008.
- [19] C. Trattner, P. Singer, D. Helic, and M. Strohmaier. Exploring the differences and similarities between hierarchical decentralized search and human navigation in information networks. In *Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies*, i-KNOW ’12, pages 14:1–14:8, New York, NY, USA, 2012. ACM.
- [20] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 2002.
- [21] R. West and J. Leskovec. Automatic versus human navigation in information networks. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media, ICWSM*, 2012.
- [22] R. West and J. Leskovec. Human Wayfinding in Information Networks. In *Proceedings of the 21th international conference on World Wide Web - WWW ’12*, New York, New York, USA, 2012. ACM Press.