The Effects of Navigation Tools on the Navigability of Web-Based Information Systems

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

Ilire Hasani-Mavriqi
Graz University of Technology
Graz, Austria
ihasani@iicm.edu

Sebastian Wilhelm
Graz University of Technology
Graz, Austria
s_wilhe@sbox.tugraz.at

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

ABSTRACT

In recent years, a number of approaches to improve the overall navigability of web-based information systems have been introduced – breadcrumbs and automatic linking algorithms represent two exemplary approaches. While evaluation of such approaches is a complex and multidimensional endeavor, involving cognitive, user-interface and other issues, the objective of this paper are more modest. In this paper, we aim to evaluate the effectiveness of different navigational aids for web information systems from a *network-theoretic perspective*. The main idea of this paper is to use decentralized search as a mechanism to evaluate navigational effectiveness. In our experiments on the largest Austrian online encyclopedia Austria-Forum, we compared the usefulness of two existing approaches to linking, i.e. breadcrumbs and automatic linking. Our results suggest that different navigational aids differ significantly with regard to their potential to support efficient navigation. The main contribution of this paper is a simulation-based framework that enriches the repertoire of tools for web engineers who are aiming to evaluate and improve navigability of web information systems. The framework allows engineers to assess a potential usefulness of various navigational tools even before expensive development or user studies are carried out.

Categories and Subject Descriptors

H.5.4 [**Information Interfaces and Presentation**]: Hypertext/ Hypermedia—*Navigation*; H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Evaluation/ Methodology*

General Terms

Theory, Experimentation, Measurement

Keywords

Navigability, Information Systems, Evaluation

1. INTRODUCTION

Recent investigations in user behavior, in particular navigational user behavior on the Web have shown interesting and to some extent surprising results. For example, in an empirical study on the so-called teleportation parameter α from PageRank, the authors measured the extent to which users follow links on a site [11] (the teleportation parameter α is a probability that the opposite will happen – it is a probability that users will teleport to a different page not by following links but by e.g. typing a new URL in the browser address bar). In the original PageRank design, the authors used $\alpha = 0.15$ [6], but the recent study found lower empirical values for the teleportation parameter. Moreover, the parameter varies between different web sites. For instance, the results show that user navigation on the web or on particular web sites, e.g. HelloMovies[1] exhibit teleportation factors that are significantly lower than the teleportation factors reported for Wikipedia users. These results are in line with our experience in operating the largest Austrian online encyclopedia called Austria-Forum [2]. Austria-Forum is a Wiki-based information system similar to Wikipedia but situated within a local Austrian context. The bounce rate (the proportion of users leaving the site after seeing only a single page) in Austria-Forum is high, about 0.6.

While [11] leave the investigation of potential causes for different user behavior to future work, the authors provide a number of possible explanations for these surprising results. In their opinion, the majority of users come to Wikipedia in search of a particular piece of information and once they arrive at Wikipedia, their information need is satisfied immediately. Therefore, users do not have any further reason, whatsoever, to click elsewhere.

Although this particular explanation seems plausible, we believe that there are further causes of this situation. One of the other causes might be a lack of rich navigational structures in sites with rather flat navigational structures such as Wikipedia. For instance, many users do not satisfy their information need with their first search query. Instead, users visit one of the first search results, follow links on that result page, backtrack, follow some other links, then in many cases refine their search, and so on [31]. With better navigational tools and aids such as overviews, links to similar documents, classifications, or indexes it might be possible to keep

[1] http://www.hellomovies.com/
[2] http://www.austria-lexikon.at

users on a given Web site, enrich their user experience and engage them on a deeper level. For example, the HelloMovies site provides overviews, categories, and other means of supporting exploratory navigation which may be one of the reasons why the site experiences high user engagement.

In this paper, we investigate this phenomenon from a network-theoretic point of view. We first present a network-theoretic framework that allows us to assess the theoretic suitability of a network for navigation. By measuring a series of structural network properties we can conclude that a particular network is navigable or non-navigable. Then, by simulating navigation in a network, we can empirically refine our measurement results – we can measure to which extent a network is navigable. Structural analysis and simulation allows us to model different navigational tools, or different algorithms for supporting navigation, and assess their (theoretical) effect on the overall navigability of the network. For example, this makes it possible to make statements about the feasibility of a particular navigational aid to support efficient navigation. Eventually, this might help answer whether richer navigational tools such as the ones offered by HelloMovies cause users to follow more links than for example the navigational aids provided by Wikipedia. Moreover, with this approach we can evaluate a particular navigation tool even *before* expensive developments or user tests are made, and thus identify the most promising approaches.

For the purposes of this study we compare two navigational tools through simulation – a standard navigational tool such as breadcrumbs and a sophisticated linking tool that automatically links similar pages according to the similarity and significance of their textual content. The first navigational tool provides a richer navigational structure and allows quick re-visiting of previously accessed pages – re-visitation browsing has been identified early as a frequent Web behavior pattern [28]. The second navigational tool is typically applied to improve horizontal connectivity of a poorly connected system. For both empirical evaluations we adopt the Austria-Forum dataset. Nevertheless, we are aware that promising approaches identified by our framework need to be evaluated in richer settings considering multiple dimensions and constraints, e.g. a semantic evaluation of the quality of linking, or execution of usability tests. In this paper however, we leave these aspects to future work.

The remainder of this paper is organized as follows. In Section 2 we introduce the evaluation framework. In Section 3 we apply the framework and investigate the navigability of the original Austria-Forum dataset. In Section 4, we present the effects of breadcrumbs on the navigability of Austria-Forum information network. In Section 5 we present the evaluation results of the navigability of Austria-Forum network enhanced with automatic links. Section 6 concludes the paper and provides a couple of ideas for the future work.

2. NAVIGABILITY EVALUATION FRAMEWORK

Research in network navigability was initiated by the famous small world experiment by Milgram [22]. In that experiment, randomly selected persons from Nebraska were asked to send a letter through their social contacts to a stockbroker in Boston. People were allowed to send the letter only to persons they knew by their first name. The striking result of the study was that letters that reached the destination needed on average about 6 hops, i.e. the population of the USA constituted a "small world". While the conclusions have been challenged [17], this experiment has attracted a great deal of interest in the research community.

Among others, two lines of research have been particularly interesting trying to gain insight into:

1. Structures that facilitate network navigability
2. Algorithms for efficient decentralized navigation, i.e. algorithms navigating with only local knowledge of the network.

In research related to (1), Kleinberg [14], [15], [16] and Watts [29] formalized these properties concluding that, from the structural point of view, a navigable network has a short path between all – or almost all – nodes in the network [16]. Formally, a navigable network has a low diameter bounded by a polynomial in $log(n)$, where n is the number of nodes in the network, and there exists a *giant component*, i.e. a strongly connected component containing almost all nodes [16].

Research related to (2) revealed that efficiently navigable networks possess additional structural properties. In particular, network nodes might be organized in a separate node hierarchy according to their similarity [19], [20], [29], or more generally to distance between nodes [1], [14], [15], [16], [29]. Then, it is possible to design efficient *decentralized navigation algorithms* [14], [15], [16]. These algorithms apply *greedy navigation* [15] in the following way: nodes use distance to select the next node on their way towards a given destination node. Thus, algorithms select the adjacent node with the smallest distance to the current destination node [14], [16]. The expected number of hops before reaching an arbitrary target node is *polylogarithmic* in n.

Recently, in [26] the authors abstract the notion of distance to a distance provided by a *hidden metric space*. Apart from instructing decentralized navigation such a metric governs also the network formation [4] – it induces several network structural properties such as node degree heterogeneity, i.e. power-law degree distributions and high clustering-coefficient [30]. The authors then connect observable emergent structural properties of a network with its navigability by defining a region of efficiently navigable networks in two dimensional space with clustering-coefficient and power-law exponent as dimensions.

In our network navigability evaluation framework, we summarize the previous research on network navigability and assess navigability in the following way:

1. **Network Generation.** We construct a network from an external source, e.g. articles from Austria-Forum and links between those articles. In this step, we can apply alternative network generators. For example, to test a navigational tool such as automatic linking of similar articles we apply automatic linking to each article and include generated links in the test network.
2. **Connectedness and Diameter.** We determine the distribution of strongly connected components in the generated network and create the hop plot. If there exists a giant component and effective diameter is bounded by $log(n)$ then we conclude that the network is navigable. The size of the giant component of a navigable network is not precisely defined in the literature. In our framework the networks with a giant component containing 80% to 90% of nodes are at a lower navigability bound, whereas the networks with the giant component with more than 90% are

definetely navigable. If the network is navigable we proceed to the next step.

3. **Degree Heterogeneity and Clustering.** We measure the degree distributions and estimate power-law exponents of these distributions. Additionally, we measure average clustering-coefficient of the network. If those values fall into the region of navigable networks as defined by [4] we conclude that network is efficiently navigable and we proceed to the next step.
4. **Greedy Navigation.** We simulate greedy navigation in the network by choosing randomly 1,000,000 node pairs (we select only low degree nodes to make the navigation task non-trivial). Each simulation begins at a start node and heads towards a destination node. As a hidden metric space we adopt an external node hierarchy, e.g. obtained by a hierarchical clustering algorithm operating on the node textual content. Then, we measure the success rate of the greedy navigation and its stretch [5], i.e. the ratio of the number of greedy hops to the global shortest path. These two measures refine the results of the third step and provide insight into the degree of the network navigability.

In previous work, we focussed on studying different aspects of tagging systems, such as navigability [13], emergent semantics [18], social classification [33] and tagging motivation[27]. In [12], we applied a similar evaluation framework to social tagging systems. However, in that paper we alternated node hierarchies and evaluated the suitability of those hierarchies to support efficient navigation. Here, we always use one and the same hierarchy but we alternate networks by incorporating linking output of different navigational tools. In that way we are able to investigate the suitability of such navigational aids for facilitating efficient navigation.

3. AUSTRIA-FORUM NAVIGABILITY

In this Section we apply our framework to evaluate the navigability of the original Austria-Forum network. Currently, Austria-Forum consists of around 45,000 textual articles with almost as many media objects. The articles are organized hierarchically into several lexicon and sub-lexicon. Lexicon are typically provided with index pages with links to included articles and/or sub-lexicon.

Network Generation. First, we generate a network representation of Austria-Forum. The network contains articles reachable from the start page by following links between articles and lexicon index pages. The network contains 46,904 nodes and 102,325 links.

Connectedness and Diameter. The calculation of the distribution of the strongly connected components reveals that the largest strongly connected component contains only around 25% of nodes (see Fig. 1). From that, we can conclude that the network is not navigable.

4. NAVIGABILITY EFFECTS OF BREADCRUMBS

Because of the limited inherent navigability of Austria-Forum, we investigate the effects of additional tools aimed at augmenting navigation in this system. We start by evaluating the effects of a popular navigational aid such as breadcrumbs.

Network Generation. We introduce breadcrumbs in Austria-Forum network by linking the complete path to the top (root) lexicon from a given article (similar to diverse Web directories, see e.g. breadcrumbs in Google Directory[3]). The newly generated network has 46,915 nodes and 260,275 links.

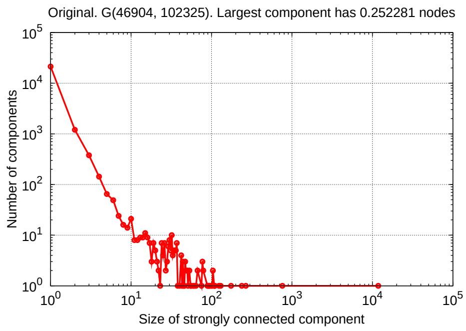

Figure 1: Distributions of strongly connected components in the original Austria-Forum network. The largest strongly connected component contains only about 25% of nodes. Therefore, the network is not navigable.

Connectedness and Diameter. Now, the network is completely connected and the largest strongly connected component has 99.99% of nodes (see Fig. 2a). The effective diameter (see Fig 2b for the complete hop plot), i.e. the length of shortest path within which 90% of node pairs are reachable from each other is 2.9 and is clearly bounded by $log(n)$. Therefore, we can conclude that this network is navigable. In the next step, we want to estimate the network's navigational efficiency.

Degree Heterogeneity and Clustering. In this step we calculate the distributions of out-degree, in-degree, and clustering-coefficient (see Fig. 3). We estimate power-law exponents γ of the degree distributions using goodness-of-fit based method [7]. We also compute the average clustering-coefficient C. According to [4] for a particular combination of the values we can conclude that a network is efficiently navigable. From our dataset, we can obtain $\gamma = 2.01$ for in-degree distribution and $\gamma = 2.36$ for out-degree distribution. These values require $C > 0.2$ and $C > 0.25$ respectively. For Austria-Forum with breadcrumbs we obtained $C = 0.86$. This allows us to conclude that the network is efficiently navigable. Next, we set up simulations to estimate the degree of the network navigability.

Greedy Navigation. In the final evaluation step we select 1,000,000 node pairs at random. The first node in the pair is the start node for a greedy decentralized navigator whereas the second node is the destination node. Both nodes are selected with $deg \leq 5$. The greedy navigator adopts a simple strategy as described in Section 2. As the background knowledge for the simulator we use lexicon hierarchy from Austria-Forum. This hierarchy is manually created by the Austria-Forum editorial board. Figure 4 shows the global success rate s and stretch τ, as well as their distributions versus observable shortest path in the network. Success rate is 100% and stretch values indicate that the greedy navigator finds its way towards the destination node in most cases in an optimal way, i.e. the number of hops h is in most cases equal to the global shortest path l. In case of longer shortest paths the greedy navigator needs more steps (e.g. $\tau = 1.13$ for $l = 4$) which means that in many cases the navigator needs to make a single additional hop as compared to the

[3]http://www.google.com/dirhp?hl=en

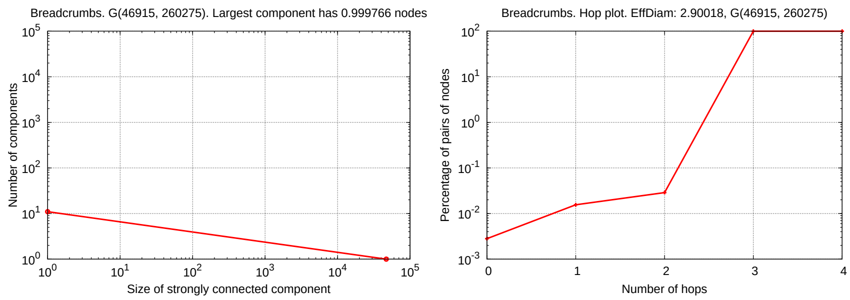

(a) Breadcrumbs AF: SCC (b) Breadcrumbs AF: Hop plot

Figure 2: Distributions of strongly connected components and hop plot in the Austria-Forum network enriched with breadcrumbs. Left: By introducing breadcrumbs, a giant component containing almost all nodes emerges. The first prerequisite for a navigable network is therefore fulfilled. Right: The effective diameter of Austria-Forum network with breadcrumbs is low. Only 2.9 hops are needed to reach 90% of nodes from each other. With 4 hops we can reach all nodes (the hop plot calculations are approximative – it is possible that small number of paths longer than 4 exist in the dataset). Thus, the second precondition is satisfied and the network is navigable.

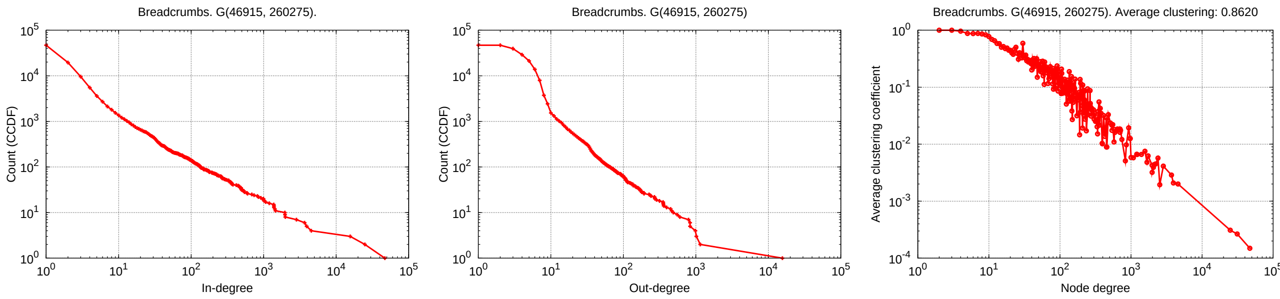

(a) Cum. In-Degree Dist. (b) Cum. Out-Degree Dist. (c) Clustering-Coefficient Dist.

Figure 3: Degree and clustering-coefficient distributions of Austria-Forum network with breadcrumbs. Left: Cumulative in-degree distribution. The estimation of the power-law exponent shows $\gamma = 2.01$. With the particular value of the average clustering-coefficient we interpret the network as efficiently navigable. Center: Cumulative out-degree distribution. The estimation produces $\gamma = 2.36$. Similarly to the in-degree power-law exponent this value is the evidence of an efficiently navigable network. Right: Clustering-coefficient distribution. The average clustering-coefficient is $C = 0.86$. Thus, the network is strongly clustered and therefore efficiently navigable.

global shortest path to find the destination node.

These results suggest that breadcrumbs seem to support efficient navigation – as breadcrumbs help organize the articles into a fully connected hierarchy of lexicon this is an expected result. However, a short analysis of the hop path structure reveals that most of the paths first go through a high degree node that is situated in the top hierarchy levels and only then to its final destination in some of the sub-nodes of the high level nodes. For example, the top node, i.e. the homepage of Austria-Forum is visited in 23% of cases, whereas the index page of the largest AEIOU lexicon is visited in 47% of navigation sessions. AEIOU lexicon contains more than 15,000 pages which are for the navigator only one hop away from the index page. In real settings users are not able to find the destination page among 15,000 other pages in a single step easily.

This is an essential limitation of our framework, which answers the question of the theoretical suitability of a particular navigational tool. Practical issues need to be evaluated by other means such as user tests.

5. EFFECTS OF AUTOMATIC LINKS ON NAVIGABILITY

5.1 Automatic Linking of Related Articles

A large fraction of articles in Austria-Forum has only few links or very often an article has no links at all. For example, AEIOU articles have been imported from the print version of the lexicon without links at all. Because of a large number of articles in that lexicon manual insertion of links by editorial team is very tedious. Moreover, many articles in AEIOU lexicon are only short descriptions of a place, person, or a town in Austria.

Nevertheless, there are articles in other lexicons that are related to those in AEIOU. For example, the AEIOU article on Bruno Kreisky is related to articles on Bruno Kreisky in the Monuments lexicon, the Post stamps lexicon, and so on. Thus, we could use semantic relatedness between such articles to enrich the system with automatically generated links and to improve, in this manner, its overall navigability. With this approach we can transform a weakly connected system to one that is at least “reasonably” connected. The final result of this action would be a system with horizontally connected articles (as opposed to vertical connectedness of e.g. the breadcrumbs network) – in essence a connectedness very similar to that of Wikipedia.

In the past, the idea of automatic links has been applied in several information systems. Generally, automatic links connect concepts from an article with other articles that describe these concepts. In particular, researchers conducted studies on automatic interlinking *within a knowledge base* [2], [9], [10], as well as studies on automatic linking of concepts to *external knowledge bases* such as Wikipedia [8], [21], [23], [24].

The literature recognizes automatic links as a tool that can enhance connectivity of information systems. For example, in [9] the authors state that links are the most important markup within an e.g. wiki system as they represent semantic relations between two concepts described within articles. Thereby, a fully connected network enables users to gain networked knowledge (as opposed to fragmented knowledge) from these knowledge bases [9].

The NNexus system [9] implements such an approach. NNexus automatically links entries of a collaborative online encyclopedia PlanetMath[4] into a semantic network of concepts using metadata of the entries. The main goal of the system is to support content authors of online encyclopedias in linking their articles with existing concepts in corpora. Researchers claim that precision and recall of automatic linking increases when a classification mechanism and a set of linking policies are applied first.

Another example of automatic linking is the Linkator [2] which combines information extraction mechanisms with Semantic Web technology to detect the terms to be linked and to annotate them semantically. Furthermore, to determine the appropriate target of the link, Linkator uses the Linked Data[5] as an external source of knowledge. Thus, terms or phrases of a web page are not linked to articles from a predefined knowledge base, but the link destination is semantically determined, on the fly, with Linked Data.

Automatic linking of arbitrary web resources to e.g. Wikipedia is also known as wikification. In [8] and [21] the authors define the concept of wikification as the task of automatically extracting the most important words and phrases from the document, and identifying for each such keyword the appropriate link to a Wikipedia article.

In their further work on the topic, Gardner and Xiong [10] identify two main problems of automatic linking of related articles: determining which terms or phrases of an article to link (known as link detection) and which articles to link to (known as link disambiguation). Recently, machine-learning approaches have shown efficient results when applied to link detection and link disambiguation problems [10], [23].

However, although researchers recognized the importance of the network connectivity as a consequence of automatic links they focused more on link quality, efficiency, and scalability [9]. Our goal is to investigate and evaluate the usefulness of this navigational aid from a network theoretic perspective. Our intention is not to invent a new automatic links algorithm – we apply standard methods for this task (Section 5.2 presents shortly our approach for automatic link generation). Rather, our intention is to *quantify the effects of such a tool on the navigability of the final information network.*

[4]http://planetmath.org/

[5]http://linkeddata.org/

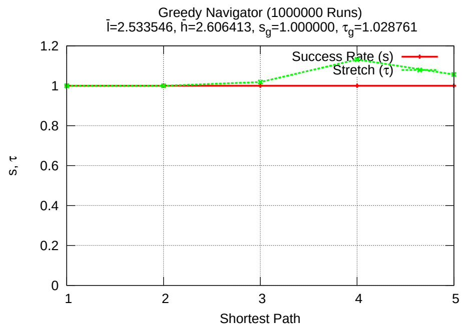

Figure 4: Success rate s and stretch τ of the greedy navigator in the Austria-Forum network with breadcrumbs. The navigator efficiency is excellent – the global success rate is 100% and the navigator needs seldom more steps than with global network knowledge. Stretch values increase slightly for longer shortest paths, e.g. $l = 4$.

5.2 Generating Automatic Links

Similarly to other automatic linking approaches our algorithm runs as a two stage process, both of which are based upon standard information retrieval methods. The first phase covers the link detection stage by focusing on identification of the article key concepts that can be used as link anchors. For link detection we apply a simplified approach of the approach described in [10] taking into account only corpus features. Thus, a standard method for calculating the key concepts is known as Term Frequency/Inverse Document Frequency – *tf.idf*. There are many variations of *tf.idf*. In our approach we applied the following formula (introduced in [25] and known as quite robust [3]):

$$tf.idf_{t,d} = (0.5 + \frac{0.5 * f_{t,d}}{max_r(f_{r,d})}) * \log \frac{N}{n_t} \tag{1}$$

Thereby, $f_{t,d}$ is the frequency of t in d, $max_r(f_{r,d})$ is the maximal frequency on any term in d, N is the total number of documents and n_t is the number of documents that contain t.

To achieve better performance we apply a couple of configurable heuristics. For example, we can configure the minimal frequency of any term $mintf$, or the idf threshold $minidf$. The performance can be further improved by applying another heuristic (can be derived from e.g. [32]) that terms with a higher idf tend to be longer, i.e. we introduce then $minwl$, wl meaning word length. In fact, we tested our algorithm with e.g. $minidf = 1.5$ and the second time with $minwl = 4$ without a significant difference in results.

The second phase deals with search for the most appropriate link destinations *within Austria-Forum knowledge base*. For that purpose we use query engine of Apache Lucene[6]. In Lucene the scoring of the documents is achieved through a mixture of boolean retrieval and the vector space model and is based also on a variation of the $tf.idf$ measure.

[6]http://lucene.apache.org/

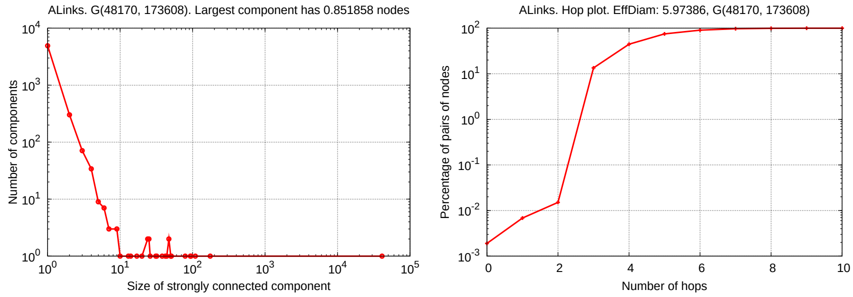

(a) ALinks AF: SCC

(b) ALinks AF: Hop plot

Figure 5: Distributions of strongly connected components and hop plot in the Austria-Forum network enriched with automatic links. Left: By automatically linking related articles a giant component containing about 85% of nodes emerges. The first prerequisite for a navigable network is fulfilled, however the network is at a lower navigability bound. Right: The effective diameter of Austria-Forum network with automatic links is bounded by $log(n)$. The second precondition is satisfied and the network is navigable.

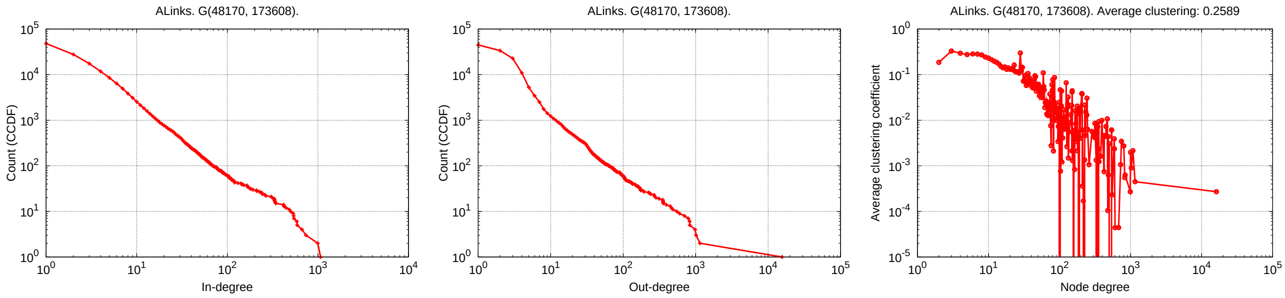

(a) Cum. In-Degree Dist.

(b) Cum. Out-Degree Dist.

(c) Clustering-Coefficient Dist.

Figure 6: Degree and clustering-coefficient distributions of Austria-Forum network with automatic links. Left: Cumulative in-degree distribution. The estimation of the power-law exponent shows $\gamma = 1.82$. This value characterizes the network as a border-line case. Center: Cumulative out-degree distribution. The estimation produces $\gamma = 1.79$. Again, the network is a border-line case. Right: Clustering-coefficient distribution. The average clustering-coefficient is $C = 0.26$ and the network is situated at a lower efficient navigation bound. Simulation should clarify the navigability degree of the network.

5.3 Navigability Evaluation

Network Generation. We apply automatic links as described in 5.2 on all articles in Austria-Forum. In order to also generate links for short articles we set the parameters as follows: $mintf = 2$ and $minidf = 1.5$. For each article we select the top three concepts to create a link. The final network has 48,170 nodes and 173,608 links.

Connectedness and Diameter. The network is connected – the largest strongly connected component has 85.19% of nodes (see Fig. 5a). However, the network is situated at the lower navigability bound. The effective diameter (see Fig 5b) is 5.97 and is bounded by $log(n)$. Therefore, we can conclude that this network is navigable and proceed to the next step to estimate its navigation efficiency. We bear in mind that the network's navigability is close to a critical point of turning to a non-navigable network.

Degree Heterogeneity and Clustering. We can now calculate the distributions of out-degree, in-degree, and clustering-coefficient (see Fig. 6) and estimate power-law exponents γ of the degree distributions. We obtain $\gamma = 1.82$ for in-degree distribution and $\gamma = 1.79$ for out-degree distribution. In [4] the authors identified lower power-law exponents, i.e. a higher degree heterogeneity as properties that support efficiency of navigation. In particular, networks with lower power-law exponents have a larger margin for clustering-coefficient values, i.e. even a network with a lower clustering-coefficient might be efficiently navigable. The lower bound for clustering-coefficient in this case is about $C > 0.2$. We obtained $C = 0.26$, which is very close to the lower bound. This is a somewhat inconclusive result – We proceed to the simulation step to gain more insight in the navigation efficiency of the network.

Greedy Navigation. As the previous analysis was inconclusive, the simulation step needs to clarify this situation. Again, we select 1,000,000 low-degree node pairs ($deg \leq 5$) at random and simulate greedy navigation with only local knowledge of the network. Success rate drops dramatically to only 3.7% (see Figure 7). In the case of success the average stretch of 1.18 is acceptable. However, such a low success rate indicates that the network is not efficiently navigable with local knowledge only. Although the existence of the giant component and a low diameter certifies that a network is navigable, this is a theoretic navigability as it holds only when a user or an agent navigating the network has the global knowledge of the network. However, for any practical application the network might be regarded as non-navigable as no user can ever have complete knowledge of such a large network.

The analysis shows that *enriching a system with automatically cre-*

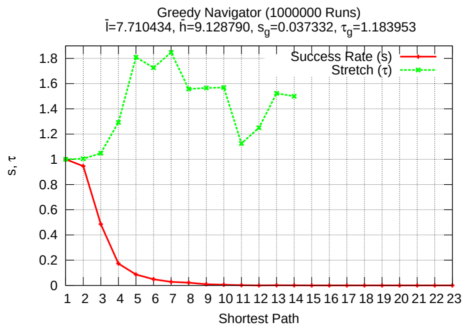

Figure 7: Success rate s and stretch τ of the greedy navigator in the Austria-Forum network with automatic links. The navigator efficiency is very poor – the global success rate is only 3.7% meaning that the network is not navigable with local knowledge only.

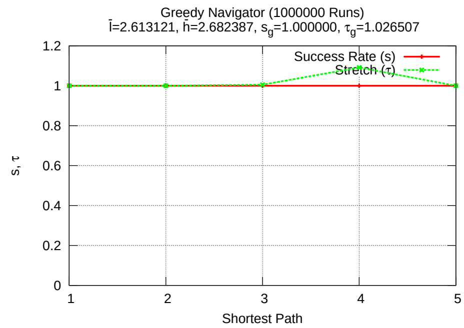

Figure 8: Success rate s and stretch τ of the greedy navigator in the Austria-Forum network with both breadcrumbs and automatic links. Because the navigator can now take shortcuts stretch values are slightly better then with breadcrumbs only.

ated links between related articles does not improve its practical navigability. The network is better connected than previously and this improves the network navigability in theory – however, for any practical application automatic links need to be combined with an additional navigation tool, e.g. breadcrumbs. Within this combination, automatic links improve horizontal connectedness of related articles whereas breadcrumbs serve as a vertical navigational aid supporting navigation with local knowledge only. Note that connectedness of Wikipedia is similar to that of Austria-Forum with automatic links, i.e. it is mainly flat and horizontal – this might be also a reason for a higher bounce rate on sites such as Wikipedia.

5.4 Combining Automatic Links and Breadcrumbs

Since the breadcrumbs network is already efficiently navigable with the combination of both navigational tools we want to achieve better horizontal connectedness of the network. In this way, the top hierarchy nodes might be relieved and more paths might take shortcuts between articles at lower hierarchy levels – this in turn might improve *practicability* of the approach because users would be less frequently confronted with huge index pages from a top level lexicon. Because we do not expect any significant differences in the first three evaluation steps we present them in a compact way without diagrams.

Network Generation. We apply both automatic links and breadcrumbs as described before. The final network has 48,229 nodes and 333,624 links.

Connectedness and Diameter. The network is completely connected and the largest strongly connected component has 99.99% of nodes. The effective diameter is 2.9. The network is navigable.

Degree Heterogeneity and Clustering. We obtain $\gamma = 1.86$ for in-degree distribution and $\gamma = 1.77$ for out-degree distribution. Clustering-coefficient $C = 0.63$ and this particular combination of parameters characterizes the network as efficiently navigable.

Greedy Navigation. Similarly to the network enriched with breadcrumbs only the success rate is 100% and stretch values indicate that the greedy navigator finds its way towards the destination node in most cases in an optimal way. In case of longer shortest paths the greedy navigator the results indicate a slight improvement in comparison with the network with breadcrumbs only (see Figure 8). The reason for this improvement is that the network has more links between articles at lower level of hierarchies, and by taking a shortcut via such a link the overall hop length is shortened. This is also indicated trough a slightly decreased fraction of paths going via the root or e.g. AEIOU index page – 21% as compared to 23% and 43% as compared to 47% respectively. This seems to be a good direction for further investigation – e.g. choosing settings for automatic links generation to optimize the improvment but not to overwhelm users with e.g. too many links.

6. CONCLUSIONS AND FUTURE WORK

The main contribution of this work is a framework for evaluating the effectiveness of navigational aids in web information systems. The idea of this paper is to use decentralized search as a mechanism to evaluate navigational effectiveness through simulations. In our experiments on the largest Austrian online encyclopedia Austria-Forum, we compared the usefulness of two existing approaches to linking, i.e. breadcrumbs and automatic linking. Our results show that interesting observations can be drawn from simulations alone, such as that breadcrumbs have useful properties and automatic linking approaches exhibit certain limitations with respect to navigability. While our empirical results are limited to the Austria-Forum dataset and the two studied approaches to linking, the introduced method for evaluation is general, i.e. it can be applied to evaluate other linking approaches in other systems.

Our results are relevant for engineers of large-scale web systems who aim to improve navigational properties of their systems and for researchers interested in evaluating navigability.

Natural expansions of our work include in the first place semantic evaluation of links. This evaluation is currently under way. The evaluation results can be in turn used as a further input to adjust the automatic links algorithm. Another interesting ideas for future work include, for example, replacing simulations with actual user click-data or studying further linking approaches in the context of other web information systems. For instance, an interesting study would be to investigate the effects of different navigational tools on the Wikipedia dataset with the goal on shedding more light on rather disappointing results of user navigational behavior on Wikipedia.

7. REFERENCES

[1] L. Adamic and E. Adar. How to search a social network. *Social Networks*, 27(3):187 – 203, 2005.

[2] S. Araujo, G.-J. Houben, and D. Schwabe. Linkator: enriching web pages by automatically adding dereferenceable semantic annotations. In *Proceedings of the 10th international conference on Web engineering*, ICWE'10, pages 355–369, Berlin, Heidelberg, 2010. Springer-Verlag.

[3] R. A. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.

[4] M. Boguñá, D. Krioukov, and K. C. Claffy. Navigability of complex networks. *Nature Physics*, 5:74–80, Jan. 2009.

[5] M. Boguñá, F. Papadopoulos, and D. Krioukov. Sustaining the Internet with hyperbolic mapping. *Nature Communications*, 1:62, Sept. 2010.

[6] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of the seventh international conference on World Wide Web 7*, WWW7, pages 107–117, Amsterdam, The Netherlands, The Netherlands, 1998. Elsevier Science Publishers B. V.

[7] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, Feb 2009.

[8] A. Csomai and R. Mihalcea. Linking documents to encyclopedic knowledge. *IEEE Intelligent Systems*, 23:34–41, September 2008.

[9] J. Gardner, A. Krowne, and L. Xiong. NNexus: An Automatic Linker for Collaborative Web-based Corpora. *IEEE Transactions on Knowledge and Data Engineering*, 21:829–839, June 2009.

[10] J. J. Gardner and L. Xiong. Automatic link detection: a sequence labeling approach. In *Proceeding of the 18th ACM conference on Information and knowledge management*, CIKM '09, pages 1701–1704, New York, NY, USA, 2009. ACM.

[11] D. F. Gleich, P. G. Constantine, A. D. Flaxman, and A. Gunawardana. Tracking the random surfer: empirically measured teleportation parameters in pagerank. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 381–390, New York, NY, USA, 2010. ACM.

[12] D. Helic, M. Strohmaier, C. Trattner, M. Muhr, and K. Lerman. Pragmatic evaluation of folksonomies. In *Proc. of the 20th International World Wide Web Conference (WWW 2011)*, New York, NY, USA, Apr. 2011. ACM.

[13] D. Helic, C. Trattner, M. Strohmaier, and K. Andrews. On the navigability of social tagging systems. In *Proc. of 2010 IEEE International Conference on Social Computing*, pages 161–168, Los Alamitos, CA, USA, 2010. IEEE Computer Society.

[14] 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.

[15] J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845, August 2000.

[16] 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.

[17] J. Kleinfeld. Could it be a big world after all? The six degrees of separation myth. *Society*, April 2002.

[18] C. Koerner, D. Benz, M. Strohmaier, A. Hotho, and G. Stumme. Stop thinking, start tagging - tag semantics emerge from collaborative verbosity. In *Proc. of the 19th International World Wide Web Conference (WWW 2010)*, Raleigh, NC, USA, Apr. 2010. ACM.

[19] E. A. Leicht, P. Holme, and M. E. J. Newman. Vertex similarity in networks. *Phys. Rev. E*, 73(2):026120, Feb 2006.

[20] F. Menczer. Growing and navigating the small world web by local content. *Proc. Natl. Acad. Sci. USA*, 99(22):14014–14019, 2002.

[21] R. Mihalcea and A. Csomai. Wikify!: linking documents to encyclopedic knowledge. In *Proceedings of the sixteenth ACM Conference on Information and Knowledge Management*, CIKM '07, pages 233–242, New York, NY, USA, 2007. ACM.

[22] S. Milgram. The small world problem. *Psychology Today*, 1:60–67, 1967.

[23] D. Milne and I. H. Witten. Learning to link with wikipedia. In *Proceeding of the 17th ACM Conference on Information and Knowledge Management*, CIKM '08, pages 509–518, New York, NY, USA, 2008. ACM.

[24] S. P. Ponzetto and R. Navigli. Knowledge-rich word sense disambiguation rivaling supervised systems. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 1522–1531, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.

[25] G. Salton and C. Buckley. Term-weighting approaches in automatic text retrieval. *Inf. Process. Manage.*, 24:513–523, August 1988.

[26] 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.

[27] M. Strohmaier, C. Koerner, and R. Kern. Why do users tag? Detecting users' motivation for tagging in social tagging systems. In *International AAAI Conference on Weblogs and Social Media (ICWSM2010), Washington, DC, USA, May 23-26*, Menlo Park, CA, USA, 2010. AAAI.

[28] L. Tauscher and S. Greenberg. Revisitation patterns in world wide web navigation. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '97, pages 399–406, New York, NY, USA, 1997. ACM.

[29] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 2002.

[30] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, June 1998.

[31] R. W. White and S. M. Drucker. Investigating behavioral variability in web search. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, pages 21–30, New York, NY, USA, 2007. ACM.

[32] G. K. Zipf. *The Psycho-Biology of Language*. MIT Press, Cambridge, Massachusetts, USA, 1935.

[33] A. Zubiaga, C. Koerner, and M. Strohmaier. Tags vs shelves: from social tagging to social classification. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pages 93–102. ACM, 2011.