

Bidirectional Fragment to Fragment Links in Wikipedia

Szymon Olewniczak, Tomasz Boiński and Julian Szymański

Gdańsk University of Technology, Poland

szymon.olewniczak@eti.pg.edu.pl

tomasz.boinski@eti.pg.edu.pl

julian.szymanski@eti.pg.edu.pl

DOI: 10.34190/EKM.20.192

Abstract: The paper presents a WikiLinks system that extends the Wikipedia linkage model with bidirectional links between fragments of the articles and overlapping links' anchors. The proposed model adopts some ideas from the research conducted in a field of nonlinear, computer-aided writing, often called a hypertext. WikiLinks may be considered as a web augmentation tool but it presents a new approach to the problem that addresses the specific nature of Wikipedia. The system, rather than working on the HTML level, works on the wiki code level and augments the Wikipedia content on the parsing phase. WikiLinks is shipped with a new kind of wiki code parser – a parallel markup parser, that allows applying externally-stored links to the content in an elegant way. The system also addresses the problem of bidirectional links' anchors adjustments on the articles' changes, provides a special link evolution procedure and handles the link's visibility problem. According to our knowledge, this is the first attempt that brings bidirectional links directly to Wikipedia. One of the possible applications of bidirectional fragment-fragment links are associative links. An associative link is a connection between two document fragments that are related to each other in some sense. The new link type extends the current Wikipedia linkage model, dominated by definitional links and provides us an additional tool for content organization, that might be useful for still-growing Wikipedia's article base. The associative links are from their nature more subjective than definitional ones. In order to test the relevance of this new linkage model, we conducted an experiment that checked if the associative links created by one person are understandable to others. The obtained results are very promising and show the usefulness of a new linkage model for Wikipedia.

Keywords: hypermedia, open hypermedia, world wide web, wiki, wikipedia, wiki parser

1. Introduction

It all began in 1945 when Vannevar Bush introduced the concept of a machine that allowed the user to create explicit connections between pages of different documents (Bush 1945). Not much time had passed when the concept became the foundation of a new scientific discipline called hypertext. The word 'hypertext' itself was created by Theodor Holm Nelson to indicate "a body of written or pictorial material interconnected in such a complex way that it could not conveniently be presented or represented on paper" (Nelson 1965). The idea was very appealing and during the years it resulted in many different systems which have formed five main hypermedia paradigms: navigational hypertext (standardized in the early '90s by the Dexter Hypertext Reference Model (Grønbaek & Trigg 1992)), stretchtext (Nelson 1987), taxonomic hypertext (Van Dyke Parunak 1991), argumentation support structures (Conklin & Begeman 1987) and spatial hypertext (Marshall et al. 1991).

Navigational hypertext is the earliest and most popular form of electronic literature. It is built around the concept of nodes, which are containers of information, and links that connect the nodes in an arbitrary way. The fragments of nodes between which the links are established are called *anchors*.

Links can take many different forms (Gunder 2002, Landow 2006). They can connect entire nodes, connect a fragment of a node with an entire node (NoteCards (Halasz et al. 1987), Hyperties (Shneiderman 1987), Guide (Brown 1987), KMS (Akscyn et al. 1988), HyperCard (Smith & Bernhardt 1988), Microcosm (Fountain et al. 1990)) or connect two (or more in the case of n-ary links) nodes' fragments (Augment/NLS (Engelbart 1962), Xanadu (Nelson 1982), Neptune (Delisle & Schwartz 1986), Intermedia (Yankelovich et al. 1988), World Wide Web (Berners-Lee et al. 1994), Devise (Grønbaek & Trigg 1994), Hyper-G (Andrews & Kappe 1994), HyperDisco (Wiil & Leggett 1996), Chimera (Anderson et al. 1997)). In addition links may also be unidirectional (Augment/NLS (Engelbart 1962), KMS (Akscyn et al. 1988), Hyperties (Shneiderman 1987), Guide (Brown 1987), HyperCard (Smith & Bernhardt 1988), Microcosm (Fountain et al. 1990), World Wide Web (Berners-Lee et al. 1994)), when they can be followed only from the source anchor or bidirectional (Xanadu (Nelson 1982), Neptune (Delisle & Schwartz 1986), Intermedia (Yankelovich et al. 1988), Devise (Grønbaek & Trigg 1994), Hyper-G (Andrews & Kappe 1994), HyperDisco (Wiil & Leggett 1996), Chimera (Anderson et al. 1997)) when they can be followed from both source and destination anchors.

Different link types may form new forms of hypertext rhetoric that allow seeing the intertextual relations on a deeper level (Atzenbeck & Nürnberg 2019). Unfortunately, the World Wide Web, the most popular hypertext system, offers only simple unidirectional links, serving as "the lowest common denominator of hypertext" (Atzenbeck & Bernstein 2018).

In this paper, a WikiLinks system is presented which introduces overlapping, bidirectional fragment-fragment links into the World Wide Web based system - Wikipedia. The new linkage mechanism allows the users to create links with greater precision and opens new perspectives for links' semantics

The paper is organized as follows: Section 2 describes the missing features of the World Wide Web, introduces the idea of a wiki, and describes the proposed extension to it. Section 3 overviews the approaches that extend the Web linkage mechanism. Section 4 outlines the WikiLinks system with its approach to the links and their versioning. Section 5 describes the idea of associative links. Section 6 reports the results of the study that tests the relevance of associative links on Wikipedia's article base. Section 7 concludes the results of the work and provides some issues for the future.

2. Problem statement

The World Wide Web is a navigational hypertext system based on a client-server model (Atzenbeck et al. 2017). The anchors in the Web are embedded, so inserting a new link requires the page content change. And although the first Web browser allowed the users to edit locally stored HTML files (Berners-Lee et al. 1994), the current browsers can only be used to viewing, not as a publishing platform¹.

Links in the World Wide Web are unidirectional. They can connect a page fragment with another page or a page fragment with some point in another page, which is explicitly defined by the destination page author through `<a id>` tag (Danilo et al. 2017).

Not much time has passed since the Web introduction when the users tried to overcome its limitations using server and client-side scripting techniques. One of the prominent examples of those attempts is the idea of a wiki, created by Ward Cunningham in 1995². The concept was to transform the World Wide Web from exploratory into editable hyperspace, where everyone can add new content easily, directly from the browser (Bo Leuf 2001). In 2001 the idea was taken on by Larry Sanger and adopted for the Wikipedia project³.

The MediaWiki engine, driving the Wikipedia, has not only facilitated the editing for the World Wide Web but also introduced some other interesting hypertext features that were lacking in the Tim Berners-Lee project. Firstly, it provides version management for the pages⁴ and permanent links to the specific page revisions⁵. Secondly, MediaWiki introduces node-fragment bidirectional links through the mechanism called backlinks⁶. Another interesting feature is page previews that allows the user to see the link's destination summary without a need for jumping⁷.

Despite the MediaWiki's improvement in linking facilities over the World Wide Web model, it still lacks some important navigational hypertext features. One of them is overlapping bidirectional fragment-fragment links between fragments of the articles. There were some approaches to introduce this type of link into the WWW but those solutions have limited use in the context of Wikipedia, where the content changes often and constantly. This raises the need for a new system that will explicitly deal with the bidirectional fragment-fragment links in the Wikipedia context.

¹ We refer here to the basic functionality of the browsers that results directly from the design of the WWW and HTTP protocol.

² <http://wiki.c2.com/?WikiHistory>

³ https://en.wikipedia.org/wiki/History_of_Wikipedia

⁴ https://en.wikipedia.org/wiki/Help:Page_history

⁵ https://en.wikipedia.org/wiki/Help:Permanent_link

⁶ https://en.wikipedia.org/wiki/Help:What_links_here

⁷ https://www.mediawiki.org/wiki/Page_Previews

2.1 Wikipedia as a knowledge management platform

It cannot be denied that Wikipedia has become one of the most important knowledge-sharing platforms on the Internet. According to Alexa rank, it is the 12th most popular Web service⁸, with over 20 billion pages views per month⁹. Wikipedia is often the important source of information for students, scholars, and knowledge workers.

The success of Wikipedia encouraged many organizations to use the wiki engines as the knowledge management platforms, often with the success (Wagner 2004). It shows that wiki's idea of a simple content sharing and editing goes beyond the encyclopedic applications and can be useful in many other cases, where knowledge management is important.

Despite its impressive information volume and social impact, Wikipedia is based on relatively simple software – MediaWiki with limited linkage capabilities. The overlapping bidirectional fragment-fragment links, as a more general form of standard web links, introduces additional precision to linkage mechanism and helps with the disambiguation of knowledge.

2.2 Technical concerns of overlapping bidirectional links

One of the reasons why the bidirectional links are not commonly available in current hypertext systems is their technical complexity. We are used to the idea, popularized by the WWW, that a link is bounded to a concrete document and contains an address that points to some resource. In the case of bidirectional links, a link is “owned” by both of its sides and the change in one of them must be reflected in the other.

The common solution to this problem is storing links separately from the content (e.g. Xanadu (Nelson 1982), Neptune (Delisle & Schwartz 1986), Intermedia (Yankelovich et al. 1988)). The additional benefit of this solution is the possibility of using multiple link sources. This allows, among others, to store public links and user's private links separately and apply them together to the same content.

The separate link storage comes with the cost of additional computation required when the content is editing since we must adjust all the link anchors on each content change. This also raises the problem of the relation between content and links versioning. All these concerns must be addressed in order to implement the proposed linkage mechanism.

2.3 The applications of overlapping bidirectional links

The extended linkage mechanism opens new possibilities both for the readers and the authors of the content. Through the overlapping anchors, the authors can create many linkage layers regarding their needs and the users can create private links reflecting their own thoughts and associations.

The bidirectional links form new types of hypertext rhetoric, that include:

- associative links – which connects fragments that relate to each other in some sense,
- narration splits – which establishes alternative paths for the current discourse,
- transclusions – which connects the common fragments that exist in different contexts.

There are definitely many more applications of the new mechanism which may reveal when implemented on a larger scale. We believe that the users, equipped with the new link type, can go far beyond our expectations and discover many other interesting use cases.

3. Related Work

Since the early days of the Web, there were many attempts to overcome its limitations by different web augmentation techniques (Bouvin 1999, 2002, Díaz & Arellano 2015). The extension of linking mechanisms in the WWW was primarily achieved through integration with various Open Hypermedia Systems. The OHSes, the third generation of hypermedia systems (Atzenbeck et al. 2017), were designed as an extension to desktop applications, that enhanced them with linking mechanisms. When integrating with the WWW, OHS creates a layer between the user and Web resources that introduce additional linkage facilities.

⁸ <https://www.alexa.com/topsites> (As of May 20, 2020)

⁹ <https://stats.wikimedia.org/#/all-projects> (As of May 20, 2020)

There are several possible architectures of the WWW-OHS integration (Anderson 1997, Bouvin 2002). Firstly, we can do it on the server level, either by enhancing an OHS server to support HTTP or joining a WWW server with the OHS client. The first solution was used by the Hyper-G server which can be accessed both from the dedicated client Harmony or using the Web browser (with some limitations) (Maurer 1996). The second idea was used in the Microcosm Distributed Link Service (Carr et al. 1995) or Chimera OHS (Anderson 1997).

Another solution is client-side integration. This technique modifies the content from the Web in the client's browser, extending it with additional links. There are two main approaches here. First is a *pre-rendering page modification* where the raw data downloaded from the Web server is processed by an augmentation tool. This solution was used by DHM/WWW (Grønbaek et al. 1997), which aimed to apply the Dexter Reference Model on top of the World Wide Web, and its successor Navette (Bouvin 1998). The alternative approach is a *post-rendering page modification* that manipulates not the raw HTML data but using its Document Object Model. This approach was taken by Webwise (Grønbaek et al. 1999) and the Arakne Framework (Bouvin 2002, Bouvin et al. 2002).

We should also mention the W3C XLink (Orchard et al. 2001) standard which extended the Web linkage model by providing many new features like external, bi-directional, typed or n-ary links. Although a serious improvement over existing mechanisms, it never got a wider adoption in real life systems.

There are also many Web annotation tools available like Diigo¹⁰, Factualnote¹¹ or Hypothes.is¹² which may be used to unidirectionally connect page fragments with some other web resources.

All presented solutions addressed the problem of missing link types in the Web but none of them focused on the specific nature of wikis with their explicit versioning and constant evolution of content. When we consider these features, introducing externally-stored links becomes a nontrivial task.

4. WikiLinks

The WikiLinks system extends Wikipedia (and other wikis that use the MediaWiki engine) with the capability of creating overlapping, bidirectional links between fragments of the articles. The links may come from many different sources. In principle, each user can create their private links on top of Wikipedia's articles. The system can be classified as the OHS that is integrated with Wikipedia on the server-side level. Clients connect to the WikiLinks server through the REST API which allows them to browse Wikipedia, create and delete links in the articles. The system is created using PHP 7 with the Symfony framework and MySQL database as a storage.

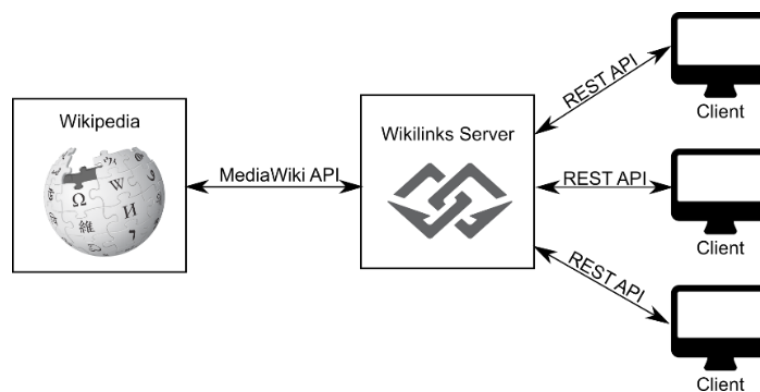


Figure 1: The architecture of WikiLinks

4.1 Links

Links in WikiLinks are first-class objects and are stored separately from the page content. Every link is composed of one or two anchors, has a type and can hold an arbitrary number of key-value pair attributes. An anchor may point to an entire page or its fragment and can be defined as a link's source or a destination. The fragment in

¹⁰ <https://www.diigo.com/>

¹¹ <http://factualnote.com/>

¹² <https://web.hypothes.is/>

WikiLinks is a continuous span of text and must fit entirely in one paragraph of an article. This limitation is imposed by the links evolution mechanism which we will discuss later.

Embedding externally-stored links in HTML documents is quite challenging, especially when links come from many different sources and their anchors are arbitrarily overlapped. This is because HTML and other XML-based formats forbid tags overlapping. For example syntax like: `This is <a>bold. and something else.` is not allowed.

There are two possible approaches to this problem. Firstly, we can create some tag splitting mechanism that will break links' anchors into several HTML tags to keep a document structure valid. Another solution is to separate the page's textual content from the HTML tags and then construct the HTML document combining tags and links' anchors. The second solution is called a *parallel markup* and was proposed by Ted Nelson as an alternative to the dominant paradigm of an *embedded markup* that can better support the hypertext structures (Nelson 1997).

In MediaWiki, the second approach may be used quite straightforwardly by replacing the standard wiki code parser with a parallel markup parser. This solution is used by us in the WikiLinks system. The Wikilink's parser takes the wiki code in input and outputs the plain text content of an article with the corresponding set of HTML tags. Each tag has information about its position in a document and its length. After the parsing, the tags set is combined with the anchoring information. Finally, all the data is sent to the users' client application which performs the rendering. The tags from the parsing process are represented as standard system links but with only a single anchor. These links are also called one-sided links or tag-links.

4.2 Versioning support

Versioning in a hypermedia environment is not an easy task both from the technical and cognitive point of view (Østerbye 1992). During the years many approaches to this problem were proposed which can be assigned into three main models: *the composition model*, *the total versioning model*, and *the product versioning model* (Nguyen et al. 2005).

MediaWiki uses *the composition model* where every atomic data container (a page and an uploaded file in MediaWiki) stores a separate version history. Because all links in MediaWiki are embedded, there is no problem in anchor updating when the source or destination page is changed.

The problem arises when we want to add an additional linkage layer on top of MediaWiki articles, and treat links as first-class objects. In this case, we must provide an additional links evolution procedure that will adjust the anchors when the article is changed.

In WikiLinks, there are two atomic data containers: the page and the link. They are both immutable objects and are identified by random UUIDs¹³. The page is a specific revision of a Wikipedia article, parsed by the specific version of the Wikilinks parser. When the parser version changes, or a new revision of a Wikipedia article is published, the new page is created. The link is composed of one or two anchors that point to the pages or their fragments. The links have also a special lineage parameter that allows identifying the following versions of the link created by the link evolution procedure. The lineage is filled with a random UUID for the newly created links or copied from the predecessor when the link is a new version of an existing one.

There are two sources of links in WikiLinks. The first are the tag-links, created from wiki code during the parsing procedure. As they come from wiki markup where syntax elements are not individually identifiable objects, they are not versioned. The second source are the user-links – bidirectional fragmentfragment links created by the system users which are evolved.

When a new revision of an article appears on the Wikipedia server or a new version of the WikiLinks parser is installed, the evolution procedure begins. In the first step, it compares the lines from the old and new page's revisions (the lines from WikiLinks parser output are used that are cleaned from a formatting syntax) using the line diff algorithm (Hunt & McIlroy 1975) and pairs the old lines with their new versions. We must remember that this is a heuristic process that may sometimes lead to errors, but since MediaWiki doesn't provide us the

¹³ Universally unique identifier in the 4th version

information about exact changes between the following revisions, we must rely on that. After the line pairing, the in-line evolution procedure starts.

The in-line evolution procedure begins by identifying the links that point to the old line. Then it runs charter diff algorithm (diff that operates on single characters, not the entire lines), comparing the old and the new line. The result of the diff is then used to adjust the new anchor scope by moving, shrinking or expanding it. After that, the old link object is cloned and the old anchor is replaced. The new link shares the lineage of its parent but has its own UUID.

As we have mentioned before links in the WikiLinks system are immutable objects. This causes that every page revision generates a bunch of new links revisions with their anchors updated. For two-sided links, it may lead to a link visibility problem (Østerbye 1992, Whitehead 2000, Soares et al. 2000, Whitehead 2001, Griffiths et al. 2003, Nguyen et al. 2005). The visibility issue appears when we have two different link versions pointing to the same anchor and must decide which to choose.

We will discuss it through an example. Imagine two Wikipedia articles *A* and *B* in which fragments are connected by *v1* link (Figure 2a). Now let's consider that the new revision of the *A* article has appeared. It launches the link evolution procedure and generates a new version of the link: *v2* (Figure 2b). Now we have two links pointing to the *B* anchor and must decide which to choose. The situation gets even more complicated when the new revision of *B* appears (Figure 2c). Which links should be visible in the successive article revisions?




Figure 2: An example of a link evolution. An empty arrow indicates a hidden anchor for the page revision

In WikiLinks, we use a strategy where links always try to point to the newest article revision. We believe that it's the most natural for users that are used to web-links. In our example the link visible in *B* will point to the anchor on *A'* and the link visible in *A'* will point to *B'*. As we can see, the proposed links evolution and visibility system leads to some asymmetries (we see different links depending on the side we are currently viewing) but on the other hand, it guarantees that the links that are visible in the last revisions always point to the newest versions of the articles.

5. Associative links

When we look carefully at the Wikipedia links semantic, we will discover they mostly represent something that can be described as a definitional link or more formally – *IS-A* relation. The link's source anchor is usually a term that is connected with an article that defines it. For example, let us consider the first sentence from the "Gdańsk University of Technology" article from the English version of Wikipedia:

The Gdańsk University of Technology (GUT; Polish: Politechnika Gdańska) is a technical university in Gdańsk-Wrzeszcz, and one of the oldest universities in Poland.

All the links in this fragment connects the terms: *Polish, technical university, Gdańsk, Wrzeszcz, Poland* with the articles that describe them.

One reason for selecting this linkage model is the lack of anchors overlapping in Wikipedia. Having many link types may lead to anchor conflict that cannot be easily solved by the tools provided by HTML. To prevent this we must decide on one linkage model and stay with that.

On the other hand, having multiple anchors over the same text fragments allows us to experiment with different linkage models and layers. In addition, the anchors can also be longer since we don't need to fear of conflict. One of the possible additional linkage models that utilizes the bidirectional fragment-fragment links and may be useful in Wikipedia, are associative links.

The associative link represents a relation between two fragments that have something in common. This connection type differs from definitional links in several ways. Firstly, both sides of associative links are some longer fragments, from one sentence to several paragraphs. Secondly, the associative link makes similar sense from both sides. Thirdly, they are more subjective and individual since they are based on a concept of an association, not a definition. The idea behind the associative link is to create connections that might be useful for a reader, and would allow him/her to see the currently read text in a wider context. This link type was also very important for the hypertext pioneer Ted Nelson, who sees in them one of hypertext's most important features (Nelson 1982).

6. Experiment

The main idea behind a WikiLinks system was to create a tool that can be used to test various linking concepts on a large corpus of Wikipedia articles. One of these concepts, which is enabled by bidirectional fragment-fragment links, are associative links that we described in the previous section.

One of the biggest problems with associative links is their subjective nature. To test the relevance of this idea in Wikipedia's article base and to test if links created by one user can be understood by others, the study was conducted.

The study included 31 participants of different ages and with a different educational background. The experiment was divided into three parts. Firstly, the idea of associative links and the WikiLinks system was presented to the participants. In the second part, the members were creating their associative links between Wikipedia's articles. Finally, the participants had to evaluate the relevance of 10 randomly selected links created by the others. During the evaluation, each connection could be assigned to one of four categories:

1. *perfect* - when the connection was understandable and both sides were well chosen;
2. *too broad* - when one side of the link was too broad which makes the association less clear;
3. *too narrow* - when one side of the link was too narrow;
4. *irrelevant* - when there was no relation between two connected fragments.

After the evaluation stage each participant was asked to calculate two metrics for the given test set:

$$Validity_{perfect} = \frac{\text{Number of links marked as perfect}}{\text{Number of all links in a test set}}$$

$$Validity_{relevant} = \frac{\text{Number of links marked as perfect, too broad or too narrow}}{\text{Number of all links in a test set}}$$

$Validity_{perfect}$ reflects the fraction of links, validated by the participant, that are well-chosen from their point of view. $Validity_{relevant}$ reflects the fraction of perfect links and links that are understandable to the participant i.e. they see the reason behind it but think that the anchors' boundaries should be chosen differently.

The metrics for each participant are presented in Table 1. The average value of $Validity_{perfect}$ evaluations between all participants was 0.53 with standard deviation $\sigma = 0.12$. The average value of all $Validity_{relevant}$ evaluations was 0.84 with standard deviation $\sigma = 0.14$. Figures 3 and 4 show the distribution of both metrics.

Table 1: The $Validity_{perfect}$ and $Validity_{relevant}$ metrics calculated by each study participant

Participant	$Validity_{perfect}$	$Validity_{relevant}$
1	0.5	0.7
2	0.6	1.0
3	0.6	0.9
4	0.4	0.8
5	0.6	0.9
6	0.6	0.8
7	0.6	1.0
8	0.5	0.8
9	0.8	0.8
10	0.4	0.7
11	0.5	0.8
12	0.4	0.9
13	0.5	0.8
14	0.7	0.9
15	0.4	0.8
16	0.6	1.0
17	0.5	0.7
18	0.4	0.6
19	0.5	1.0
20	0.6	0.8
21	0.6	1.0
22	0.7	0.9
23	0.6	0.8
24	0.3	0.7
25	0.7	1.0
26	0.4	0.6
27	0.7	0.8
28	0.5	1.0
29	0.6	1.0
30	0.3	0.5
31	0.4	0.9




Figure 2: Distribution of the $Validity_{perfect}$ metric

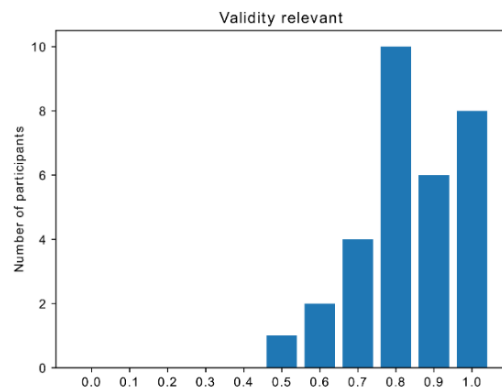


Figure 3: Distribution of the $Validity_{relevant}$ metric

7. Conclusions and future work

The WikiLinks system proved to be valid and useful. The experiment shows that the associative links created by one user are mostly understandable by the others and can provide an interesting extension to Wikipedia's linkage model. We believe that the presented solution is a good starting point for the discussion about new methods of organizing the still-growing Wikipedia information base.

The solution worked very well as a research tool. However, the system is at an early stage and there are many places where it can be improved.

Firstly, we can improve the linkage mechanism itself. One of the most important challenges here is the introduction of n-ary links to the system. However, how to implement this feature on the user interface level to make it clear and useful is still an open question.

Another challenge regarding the linking mechanisms is the improvement of the links' evolution procedure. Currently, we use a heuristic based on the diff algorithm which may generate bad results for some cases. We should explore the possibilities of making the process more reliable, either by enhancing the current solution or by trying different algorithms.

Secondly, there is still much to do in the parallel wiki code parser. The present version cleans the majority of the Wikipedia syntax without creating appropriate tag-links for it. In order to make the solution useful for the greater audience, the parallel parser must become much more advanced.

Thirdly, we should constantly look for new use cases for the proposed linkage mechanism. We are planning to conduct several other experiments that will test the relevance of navigational splits and transclusions in the context of Wikipedia.

Finally, the overall solution must be tested in the context of user experience. We want to check in which cases new types of links facilitate knowledge acquisition and usage. Additionally, we will test different link presentation methods in the context of its clarity and usefulness. We believe that there are still many unexplored applications of the bidirectional fragment-fragment links that can make a Wikipedia experience even better.

References

- Akscyn, R. M., McCracken, D. L., & Yoder, E. A. (1988). KMS: A Distributed Hypermedia System for Managing Knowledge in Organizations. *Commun. ACM*, 31 (7), pp 820–835.
- Anderson, K., Taylor, R., & Whitehead, E. (1997). Chimera: Hypertext for heterogeneous software environments.
- Anderson, K. M. (1997). Integrating Open Hypermedia Systems with the World Wide Web. *Proceedings of the Eighth ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 157–166
- Andrews, K., & Kappe, F. (1994). Soaring through hyperspace: A snapshot of Hyper-G and its Harmony client. In W. Herzner & F. Kappe (Eds.), *Multimedia/Hypermedia in Open Distributed Environments*. Vienna: Springer Vienna. pp 181–191
- Atzenbeck, C., & Bernstein, M. (2018). Interview with Andy Van Dam. *SIGWEB Newsl.* (Winter), 1:1–1:7.
- Atzenbeck, C., & Nürnberg, P. J. (2019). Hypertext As Method. *Proceedings of the 30th ACM Conference on Hypertext and Social Media*. New York, NY, USA: ACM. pp 29–38
- Atzenbeck, C., Schedel, T., Tzagarakis, M., Roßner, D., & Mages, L. (2017). Revisiting Hypertext Infrastructure. *Proceedings of the 28th ACM Conference on Hypertext and Social Media*. New York, NY, USA: ACM. pp 35–44.
- Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H. F., & Secret, A. (1994). The World-Wide Web. *Commun. ACM*, 37 (8), pp 76–82.
- Bo Leuf, W. C. (2001). *The Wiki Way: Quick Collaboration on the Web*. Addison-Wesley Professional.
- Bouvin, N. O. (1998). Designing Open Hypermedia Applets: Experiences and Prospects. *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*. New York, NY, USA: ACM. pp 281–282
- Bouvin, N. O. (1999). Unifying strategies for web augmentation. *Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia*. New York, NY, USA: ACM. pp 91–100
- Bouvin, N. O. (2002). Augmenting the web through open hypermedia. *New Review of Hypermedia and Multimedia*, 8 (1), pp 3-25.
- Bouvin, N. O., Zellweger, P. T., Grønbaek, K., & Mackinlay, J. D. (2002). Fluid Annotations Through Open Hypermedia: Using and extending emerging web standards. *Proceedings of the 11th International Conference on World Wide Web*. New York, NY, USA: ACM. pp 160–171
- Brown, P. J. (1987). Turning Ideas into Products: The Guide System. *Proceedings of the ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 33–40
- Bush, V. (1945, 7). As we may think. *The Atlantic Monthly*, 176 (1), pp 101–108.

- Carr, L., De Roure, D., Hall, W., & Hill, G. (1995). The distributed link service: A tool for publishers, authors and readers. *World Wide Web Journal*, 1 (1), pp 647–656.
- Conklin, J., & Begeman, M. L. (1987). gIBIS: A Hypertext Tool for Team Design Deliberation. *Proceedings of the ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 247–251
- Danilo, A., Faulkner, S., Moon, S., Eicholz, A., & Leithead, T. (2017). HTML 5.2 (W3C Recommendation). W3C.
- Delisle, N., & Schwartz, M. (1986). Neptune: A Hypertext System for CAD Applications. *Proceedings of the 1986 ACM SIGMOD International Conference on Management of Data*. New York, NY, USA: ACM. pp 132–143
- Díaz, O., & Arellano, C. (2015, May). The Augmented Web: Rationales, Opportunities, and Challenges on Browser-Side Tanscoding. *ACM Trans. Web*, 9 (2), pp 8:1–8:30.
- Engelbart, D. C. (1962). Augmenting Human Intellect: A Conceptual Framework (Summary Report No. AFOSR-3233). Stanford Research Institute.
- Fountain, A., Hall, W., Heath, I., & Davis, H. (1990). Microcosm: An Open Model for Hypermedia with Dynamic Linking. In A. Rizk, N. Streitz, & J. Andre (Eds.), *Hypertext: Concepts, Systems and Applications, Proceedings of ECHT'90*, Paris. pp 298–311.
- Griffiths, J., Millard, D. E., Davis, H., Michaelides, D. T., & Weal, M. J. (2003). Reconciling Versioning and Context in Hypermedia Structure Servers. In P. J. Nürnberg (Ed.), *Metainformatics*. Berlin, Heidelberg: Springer Berlin Heidelberg. pp 118–131
- Grønbaek, K., Bouvin, N. O., & Sloth, L. (1997). Designing Dexter-Based Hypermedia Services for the World Wide Web. *Hypertext 97, The Eighth ACM Conference on Hypertext*. University of Southampton, UK. pp 146–156.
- Grønbaek, K., & Trigg, R. H. (1992). Design Issues for a Dexter-based Hypermedia System. *Proceedings of the ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 191–200
- Grønbaek, K., & Trigg, R. H. (1994). Design Issues for a Dexter-based Hypermedia System. *Commun. ACM*, 37 (2). pp 40–49.
- Grønbaek, K., Sloth, L., & Ørbæk, P. (1999). Webwise: browser and proxy support for open hypermedia structuring mechanisms on the World Wide Web. *Computer Networks*, 31 (11), pp 1331 – 1345.
- Gunder, A. (2002). Aspects of Linkology: A method for the description of links and linking.
- Halasz, F. G., Moran, T. P., & Trigg, R. H. (1987). Notecards in a Nutshell. *Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics Interface*. New York, NY, USA: ACM. pp 45–52
- Hunt, J. W., & McIlroy, M. D. (1975). An algorithm for differential file comparison.
- Landow, G. P. (2006). *HHypertext 3.0: Critical Theory and New Media in an Era of Globalization (Parallax: Re-visions of Culture and Society)*. Johns Hopkins University Press.
- Marshall, C. C., Halasz, F. G., Rogers, R. A., & Janssen, W. C., Jr. (1991). Aquanet: A Hypertext Tool to Hold Your Knowledge in Place. *Proceedings of the Third Annual ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 261–275
- Maurer, H. A. (1996). *Hyper-G now Hyperwave: the next generation Web solution*. England: Addison-Wesley Longman Ltd.
- Nelson, T. H. (1965). Complex Information Processing: A File Structure for the Complex, the Changing and the Indeterminate. *Proceedings of the 20th National Conference*. New York, NY, USA: ACM. pp 84–100
- Nelson, T. H. (1982). *Literary Machines*. Mindful Press.
- Nelson, T. H. (1987). *Computer Lib/Dream Machines, Revised Edition*. Tempus Books.
- Nelson, T. H. (1997). Embedded Markup Considered Harmful. *World Wide Web Journal*, 2, pp 129–134.
- Nguyen, T. N., Thao, C., & Munson, E. V. (2005). On Product Versioning for Hypertexts. *Proceedings of the 12th International Workshop on Software Configuration Management*. New York, NY, USA: ACM. pp 113–132
- Orchard, D., Maler, E., & DeRose, S. (2001, June). XML linking language (XLink) version 1.0 (W3C Recommendation). W3C.
- Østerbye, K. (1992). Structural and Cognitive Problems in Providing Version Control for Hypertext *Proceedings of the ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 33–42
- Shneiderman, B. (1987). User Interface Design for the Hyperties Electronic Encyclopedia (Panel Session). *Proceedings of the ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 189–194
- Smith, T., & Bernhardt, S. (1988). Expectations and Experiences with HyperCard: A Pilot Study. *Proceedings of the 6th Annual International Conference on Systems Documentation*. New York, NY, USA: ACM. pp 47–56
- Soares, L. F. G., Rodrigues, R. F., & Muchaluat Saade, D. C. (2000). Modeling, authoring and formatting hypermedia documents in the HyperProp system. *Multimedia Systems*, 8 (2), pp 118–134.
- Van Dyke Parunak, H. (1991). Don'T Link Me in: Set Based Hypermedia for Taxonomic Reasoning. *Proceedings of the Third Annual ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 233–242
- Wagner, C. (2004). Wiki: A technology for conversational knowledge management and group collaboration. *Communications of the association for information systems*, 13(1), 19.
- Whitehead, E. J., Jr. (2000). *Analysis of the Hypertext Versioning Domain* (Unpublished doctoral dissertation). University of California, Irvine. (AAI9980917)
- Whitehead, E. J., Jr. (2001). Design Spaces for Link and Structure Versioning. *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia*. New York, NY, USA: ACM. pp 195–204
- Wiil, U. K., & Leggett, J. J. (1996). The HyperDisco Approach to Open Hypermedia Systems. *Proceedings of the the Seventh ACM Conference on Hypertext*. New York, NY, USA: ACM. pp 140–148
- Yankelovich, N., Haan, B. J., Meyrowitz, N. K., & Drucker, S. M. (1988). Intermedia: the concept and the construction of a seamless information environment. *Computer*, 21 (1). pp 81-96.