

A Quantitative Comparison of Semantic Web Page Segmentation Approaches

Robert Kreuzer, Jurriaan Hage^(✉), and Ad Feelders

Department of Information and Computing Sciences, Utrecht University,
Utrecht, The Netherlands

robert.kreuzer@gmail.com, {J.Hage,A.J.Feelders}@uu.nl

Abstract. We compare three known semantic web page segmentation algorithms, each serving as an example of a particular approach to the problem, and one self-developed algorithm, WebTerrain, that combines two of the approaches. We compare the performance of the four algorithms for a large benchmark of modern websites we have constructed, examining each algorithm for a total of eight configurations. We found that all algorithms performed better on random pages on average than on popular pages, and results are better when running the algorithms on the HTML obtained from the DOM rather than on the plain HTML. Overall there is much room for improvement as we find the best average F-score to be 0.49, indicating that for modern websites currently available algorithms are not yet of practical use.

1 Introduction

Web page segmentation is the process of taking a web page, and partitioning it into so-called *semantic blocks* (or segments), that we define as

A contiguous HTML fragment which renders as a graphically consistent block and whose content belongs together semantically.

Semantic blocks can in principle be deeply nested, although in practice people rarely consider nesting more than two or three levels. In this paper we do not consider the process of *labeling*, which takes a partitioning of a webpage into semantic blocks, and then assign labels to them taken from some ontology.

Human beings are very good at partitioning: even if a website is in a language we are not familiar with, it is clear to us what is an advertisement, what is a menu, and so on. Web page segmentation algorithms seek to automate this process: among the applications for automated segmentation we find mobile web, voice web, web page phishing detection, duplicate deletion, information retrieval, image retrieval, information extraction, user interest detection, visual quality evaluation, web page clustering, caching, archiving, semantic annotation and web accessibility [18].

In this paper we perform an empirical study of the usefulness of three web page segmentation *approaches*: the DOM-based approach, the visual approach

and the text-based approach. We also consider whether combining two of these approaches gives superior results. We compare the approaches by taking a representative algorithm for each, and evaluate their effectiveness in web page segmentation by comparing their outputs to a manually constructed ground truth. This should answer our first research question, which reads

RQ1: How well do different web page segmentation approaches perform on a large collection of modern websites? In particular, can we observe that the algorithms indeed fare worse on modern websites?

An additional, related research question we considered is the following:

RQ2: Can the results be improved by combining different approaches?

As we shall see in this paper, the answer to both these questions is affirmative. Although the algorithm that combines two of the approaches often outperforms the other algorithms, its effectiveness is still much below what we would like to see. A conclusion of this paper is then also that new, better algorithms should be sought to more effectively perform web page segmentation.

We have constructed two datasets: one with a number of popular websites, that we expected to be complex and therefore hard to deal with, and a second set with randomly selected websites to avoid having our results biased too much towards complicated websites. Our datasets are part of the contribution this paper offers, and are open for anyone to use¹.

The paper is structured as follows: in Section 2 we discuss the various approaches to web page segmentation, and our choice of algorithms for the comparison. We had to construct our own dataset of marked up web sites, and we discuss the process of doing so in Section 3. In Section 4 we provide the outcomes of our comparison, upon which we reflect in Section 5. Section 6 discusses related work, and Section 7 concludes.

2 Approaches and Algorithms

There are many algorithms for web page segmentation, and we cannot hope to consider them all. Instead, we opted to see which of the *approaches* to web page segmentation work well on modern websites. For each approach we chose a representative algorithm (typically the one that is generally most cited or compared against), obtained an implementation from others, or, in the absence of such an implementation, implemented the algorithm ourselves. The three approaches we consider are the DOM-based approach, the visual approach and the text-based approach.

2.1 The DOM-Based Approach (PageSegmenter)

In a DOM-based approach one simply looks at the DOM, the tree built by parsing the HTML, for cues on how to segment a page. (This tree does not include the

¹ <https://github.com/rkrzr/dataset-random>, <https://github.com/rkrzr/dataset-popular>, constructed by Kreuzer, El-Lari, Van Nuenen, and Hospes

properties added by external CSS files; these will only be present in the render tree.) The idea is that the HTML structure should reflect the semantics of the page. The quality of these approaches thus depends on how strongly this is the case. To do the segmentation they rely on detecting recurring patterns, such as lists and tables, and on heuristics, like headline tags working as separators, and links being part of the surrounding text. The approach is easy to implement and efficient to run, since one only needs to parse the HTML and not render the page. Complications are that there are many different ways to build an HTML document structure for the same content, styling and layout information is disregarded by design, and it will not work immediately in the presence of Javascript (unless you serialize the DOM first).

The PageSegmenter algorithm is an example of a DOM-based segmentation algorithm [16]. They tested their work experimentally against the TAP knowledge base [10] (which was not available anymore for our study) and on a home-made dataset consisting of CS department websites. In [17], the authors also rate the individual blocks by learning a statistical predictor of segment content quality and use those ratings to improve search results.

The main idea of the PageSegmenter algorithm is that the root-to-leaf paths of leaf nodes can be analyzed for similarities to find nodes which likely belong to the same semantic block. An example of such a path is `/html/body/p/ul/li`. Now, if multiple adjacent siblings have the same such path it is a pretty safe assumption that they also semantically belong together, as they are structurally part of the same list.

In their paper the authors mention that they exclude text nodes from the tree where the text contains *modal verbs* (such as could, should, would...) in order to decrease noise. Our implementation does not do so, because this would make the algorithm language-specific, which is something we want to avoid. More details on how we implemented the algorithm precisely can be found in [15].

2.2 The Visual Approach (VIPS)

Visual approaches most resemble how a human segments a page, i.e. they operate on the rendered page itself as seen in a browser. They thus have the most information available, but are also computationally the most expensive because pages must be rendered. They often divide the page into separators, such as lines, white-space and images, and content and build a content-structure out of this information. They can take visual features such as background color, styling, layout, font size and type and location on the page into account. To render the page we need access to a browser engine, which complicates the implementation of an algorithm. Moreover, it clearly requires external resources such as CSS files and images in order to work correctly.

For this approach we elected the VIPS (Vision-based Page Segmentation) algorithm [5,6], which appears to be the most popular web page segmentation algorithm. As indicated by the name this algorithm is based on the rendered representation of a page. It analyzes the DOM after all the styling information from CSS rules have been applied and after Javascript files were executed (and

potentially modified the tree). It is tightly integrated with a browser rendering engine since it needs to query for information such as the dimensions on screen of a given element. One thus has to decide on a fixed viewport size in advance on which the page should be rendered. Concretely, the algorithm builds a vision-based content structure, independent of the underlying HTML document, deciding during a top-down traversal whether something represents a visual block, or whether it should be subdivided further by using a number of heuristics, such as “if a sub-tree contains separators like the `<hr>` tag, subdivide”. The authors tested their algorithm experimentally by sampling 140 pages from different categories of the Yahoo directory and running their algorithm on them and then manually assessing whether the segmentation was “Perfect”, “Satisfactory” or “Failed”. We used an existing implementation from Tomas Popela²; the original implementation was not available anymore.

2.3 The Text-Based Approach (BlockFusion)

The *text-based* approach differs from the other two in that it does not take the tree structure of the HTML into account at all. Algorithms only look at the (textual) content and analyze certain textual features like e.g. the text-density or the link-density of parts of a page. These techniques are grounded in results from quantitative linguistics which indicate that, statistically, text blocks with similar features are likely to belong together and can thus be merged in a single block. The optimal similarity threshold depends on the wanted granularity and needs to be determined experimentally. The algorithms tend to be fast and easy to implement since they work independently from the DOM, but like the DOM-based approach will not work with Javascript (unless you serialize the DOM first), do not take structural and visual clues into account, and the extraction of sub-blocks requires local changes to the text-density threshold (since we can't employ document structure).

The representative chosen for the text-based approach is the BlockFusion algorithm [13]. The algorithm is based on the observation that the so-called *token density* can be a valuable heuristic to segment text documents. The token density of a text can simply be calculated by taking the number of words in the text and dividing it by the number of lines, where a line is capped to 80 characters. An HTML document is then first preprocessed into a list of atomic text blocks, by splitting on so-called separating gaps, which are HTML tags other than the `<a>` tag. For each atomic block the token density can then be computed. A merge strategy is then employed to merge blocks into progressively larger ones, if the difference between token densities of two adjacent blocks is below a certain threshold value. This is done repeatedly until no more blocks can be merged. Due to this design the algorithm does not support multiple levels of blocks by default, but by an extension in which we locally introduce a second smaller threshold value, and then call the BlockFusion algorithm on each (already merged) block, we can achieve a two-level hierarchy. BlockFusion was

² <https://github.com/topopela/vips.java>

evaluated experimentally using a dataset consisting of 111 pages. We based our own implementation of this algorithm on an open source library from one of the BlockFusion authors [12], allowing us to stay close to the original implementation. We used a text density threshold value of $\vartheta_{max} = 0.38$, which was found to be the optimal value in the experimental evaluation [13].

2.4 A Combined Approach (WebTerrain)

The WebTerrain algorithm was developed as our own contribution to the segmentation problem. The main idea was to see if we can combine the different approaches from the other algorithms in order to improve upon the end result. The algorithm is based on a novel heuristic which inspired the name of the algorithm: Firefox has a little known feature which allows the user to see a 3D-rendered version of any website (choose Inspect Element after a right-click on any given page, and then click on the cube icon). The result looks similar to a geographic terrain map. This feature works by assigning an additional depth-value to each visible element on top of the common width- and height-values, which are already used in the normal 2D-representation of the page. The depth-value is simply the tree-level of the element. Experiments with this feature revealed that the elevation profile corresponds pretty well to what we would consider the semantic blocks of a web page. The heuristic that this observation leads to has the interesting property that it combines a plain structural approach with a rendering-based approach into one, since it not only takes the DOM tree into account but also the visibility and dimensions of each element. It is not possible to tell by simply looking at the original HTML document how it will ultimately be rendered. One does not know how much space each child of an element will take up on the screen, or if it will be visible at all. For this, one needs to actually render the page.

Further details about the various algorithms, implementations, complications, and how we extracted the necessary information from the outputs of the implementations are omitted for reasons of space, but can be found in [15].

3 The Datasets

Since web page segmentation is a fairly well-studied problem, many authors have done an empirical evaluation of their algorithms. The datasets and methods used for this evaluation vary widely. There appears to be no standard dataset for this problem, instead everyone seems to create their own dataset by first randomly sampling web pages (sometimes taken from a directory site such as <http://dmoz.org>) and then manually marking up the semantic blocks and often also labeling the blocks according to a predefined ontology. To further illustrate this, we consider how the three chosen known algorithms were validated by the authors that proposed them.

For the VIPS algorithm the authors did not first create a dataset with a ground truth. Instead they ran VIPS on their sample pages and then manually

graded whether the segmentation of that page was “perfect”, “satisfactory” or “failed”. This approach is problematic on two levels: First, there is the obvious conflict of interest, since the authors themselves are grading the results of their own algorithm. Second, whether a segmentation is “perfect” or “satisfactory” is rather subjective and can thus not be repeated by others.

For the BlockFusion algorithm the authors did not use precision and recall, but instead they used two cluster correlation metrics, namely *Adjusted Rand Index* and *Normalized Mutual Information* to quantify the accuracy of the segmentation. They did first create a ground truth manually, but it is unclear whether this was done by the authors themselves or by volunteers.

For the PageSegmenter algorithm the authors did use precision, recall and F-Score in their evaluation. Differently from us they did not do this for all blocks in general on a page, but they divided the blocks into three classes first (which they call Concept, Attribute and Value) and applied the metrics to each of these classes. This again prevents a direct comparison as this division into three classes is specific to their algorithm and not applicable to other segmentation algorithms.

Before building our own dataset we investigated the datasets used by other authors to find out how they chose their sample pages, sample sizes and whether they downloaded only the HTML documents themselves, or the referenced external resources as well. Furthermore we wanted to see whether any of these datasets would be suitable for our study as well.

We found five datasets. The manually labeled ones vary in size from 105 to 515, with the exception of the TAP knowledge base [10] at a size of 9,068 which was a semantically labeled database that was used as a test-bed for the Semantic Web but is unfortunately not available anymore. The Web pages are sampled completely at random in [7], in [13] they are taken from the Webspam UK-2007 dataset [3, 4, 8] comprising over 100 million pages, which is focused on labeling hosts into spam/nospam, in [14] they first downloaded 16,000 random pages from the directory site www.dmoz.org and randomly selected sample pages. Also, [16] distinguishes between template-driven and non-template-driven Web pages (i.e. pages generated by a web page authoring system and hand-coded ones).

The block assignment was sometimes done by the authors and sometimes by volunteers, the latter being preferable to avoid bias. It is not always mentioned what granularity of blocks was used (i.e. whether only top-level blocks were marked up or sub-blocks as well), but no one specifically mentioned marking up sub-blocks which leads us to the assumption that no sub-blocks were highlighted. Since none of these datasets are available online or from the authors directly we were unable to confirm this.

One other notable observation is that all datasets seem to consist only of HTML documents without any of the external resources referenced from the pages. While this is certainly partly due to the datasets being already up to 11 years old, when web pages on average were still a lot less complex than they are now, this is not realistic anymore for websites that are currently on-line. We discuss why later on in this section.

Neither of these potential benchmarks turned out to be suitable: either the authors of the papers never replied to our inquiries, and in the one case that we did find the dataset it was not suitable for our purposes since it was focused on news websites, and only included the original HTML sources and no external resources. The latter is a problem in our case, because all algorithms that depend on a rendered representation of the page will deliver very different results for a page with all external resources and one without. In conclusion, we decided to construct our own benchmark set, which we have made publicly available (see the footnote in Section 1).

External Resources in Web Pages. Modern web pages can and typically do reference a number of external resources that are necessary to render the page properly. These include images, videos, CSS files, Javascript files, and lesser known ones like font files, JSON or XML files, favicons, and vector graphic files.

Clearly, if you render only the HTML document itself then the result will usually be vastly different from the page including all resources. Mainly for this reason we decided to build a dataset consisting of HTML documents together with all their external resources (and all links rewritten accordingly so that they point to the right files). Javascript poses a real challenge for any kind of Web page analysis. Since a Javascript program can modify the DOM arbitrarily and furthermore load in more data or other Javascript programs from external sources, it is possible that the original HTML document and the rendered page have virtually nothing in common.

In effect, if we want to be able to validate algorithms that employ the visual approach, we must include in our dataset all the external resources that a given web page needs to be visualized. To retrieve, for a given web page, all such external resources is not a trivial exercise. For example, CSS files have the ability to refer to other CSS files, and these may again refer to ever more such files. Javascript is particularly hard to deal with since it is a Turing complete language (unlike HTML and CSS). In practice we solve the problem by simply running the Javascript programs, but whether we have in fact retrieved all resources that we shall ever need is undecidable in general. The best practical solution we could find is the `wget` utility, using finely tuned parameters³. It handles all of the difficulties mentioned above except references from within Javascript programs, and it also rewrites the links so that they all point to the right locations. We found that the downloaded pages rendered identical or nearly identical to the online version in most cases. The pages that used Javascript references to an extent that they could not be properly rendered offline were excluded from the dataset (18 out of 100 for the random dataset and 30 out of 100 for the popular dataset).

Having retrieved the web pages and the associated resources, they had to be marked up so that we would have a ground truth to compare the algorithms against. We implemented a small program, called Builder, in Javascript that,

³ The magic incantation is: `wget -U user_agent -E -H -k -K -p -x -P folder_name -e robots=off the_url`

given the url of the web page, allows one of our volunteers to easily mark up the web page and store the results. It is run by clicking a bookmarklet which will load and run the Javascript program from our server. It then works by highlighting the DOM node the user is hovering over with the mouse, allowing her to select that block and showing her a menu where she can choose what type of block it is. The possible choices to classify a block were:

High-level-blocks Header, Footer, Content, Sidebar

Sub-level-blocks Logo, Menu, Title, Link-list, Table, Comments, Ad, Image, Video, Article, Searchbar

This block ontology was chosen with the goal of being comprehensive and it was divided into High-level blocks and Sub-level blocks (or level 1 and level 2 blocks) since Web pages can be segmented on different levels of granularity. E.g. a content-block can have a title, an image and an article as sub-blocks. While in principle there is no upper limit to how many levels of granularity you can have on a page, we found two levels to be sufficient in the majority of cases and have thus restricted ourselves to that.

For robustness we implemented the following solution to marking up the web pages in which the client (for marking up the web pages) and the server (serving the local version of the web pages) reside on the same machine: we first make the full page available offline using `wget`, then open that page in a browser, load the Builder, and add all the blocks and finally serialize the changed DOM to disk again. If one subsequently wants to get out all the blocks of a page one can do so using a simple Xpath query⁴.

We built two different datasets, one containing only popular pages and one containing randomly selected pages. This was done to see if the algorithms performed differently on random and on popular pages on average. For the popular dataset we took the top 10 pages from the 10 top-level categories from the directory site <http://dir.yahoo.com/>. The chosen categories were Arts & Humanities, Business & Economy, Computer & Internet, Entertainment, Government, Health, News & Media, Science and Social Science. We believe this gives us a representative sample of popular websites, although not of websites in general. We manually checked all websites whether they still rendered properly after having been downloaded and removed the ones that were broken, which left us with a total of seventy popular pages in the dataset.

For the random websites we made use of the web service from

<http://www.whatsmyip.org/random-website-machine/>

to generate a hundred links, which we then downloaded. The service boasts over four million pages in its database and the only filtering done is for adult content, which makes it sufficiently representative for the Internet as a whole. After removing pages that did not render properly offline we ended up with a random dataset consisting of 82 pages.

The marking up of the semantic blocks on these pages was done by three volunteers. They were instructed to first mark up all the level-1 blocks they could

⁴ Xpath query to get all blocks: `'//*[@data-block]'`

find and subsequently all the level-2 blocks within each level-1 block, according to the generic ontology we gave earlier in this section.

When establishing the ground truth with the help of outsiders, we must in some way be reasonably sure that this “truth” is objective enough. In other words, can we expect the intuitive understanding of what is a semantic block among our test persons to be aligned, and, moreover, in line with what the average person surfing the Web would say? The work of [16] reports an overlap of 87.5% between eight test subjects who they asked to mark up the same web pages. Although the sample is small, it corresponds to our own anecdotal experience. However, we did find that we had to be specific about the level of granularity (e.g. “the most high-level (i.e. biggest) blocks and their most high-level sub-blocks”), since there can be many levels.

4 Results

In this section we present the results of our evaluation of the four different segmentation algorithms. We tested all algorithms in a number of different configurations using a custom testing framework that abstracted the differences in implementations and normalized the results, which could subsequently be analyzed statistically [15].

First, we tested them on the two different datasets which we created for this purpose: the randomly selected dataset and the popular dataset. The first one consists of 82 random pages and the second one of 70 popular pages, all marked up by our assessors. We chose these two types of datasets to test whether the algorithms perform differently on random and on popular pages on average.

As a second variable we ran the algorithms on both the original HTML, i.e. the HTML document downloaded from the source URL via a single GET request, and the DOM HTML, i.e. the HTML document obtained by waiting for all external resources to load and then serializing the DOM. As there appears to be a trend to build pages dynamically on the client-side using Javascript, we were interested to see whether our results would reflect this. It is also of note that our tool to mark up blocks manually was browser-based and thus operated on the DOM, making the DOM HTML the true basis of our ground truth. We believe this is a more sensible basis than the original HTML, since it is what the user ultimately sees when viewing a page, and it also is what the creator of the page intended as the final result.

Finally we used two metrics to compare the generated results to the ground truth, the exact match metric and the fuzzy match metric. Both of them compare the string contents of the blocks to each other. Each block is serialized to only text with all HTML tags removed and white-space and newlines removed as well. For the exact match metric it then simply checks for string equality. This is of course a very strict criterion, as a minimally different string would be counted as false, while for most applications it would likely be perfectly sufficient. For this reason we also do a fuzzy string comparison to check for a similarity ratio of better than 0.8 between strings.

Table 1. Results for Random-HTML-Exact (left) and Random-HTML-Fuzzy (right). Note that the reported numbers are averages over 82 web pages. E.g., the reported F-score was computed by first computing the F-score for each web page, and then taking the average of these 82 F-scores. Therefore the reported values for Precision, Recall and F-score may not satisfy the formula for the F-score as given in the text.

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.03	0.06	0.04	25.99	0.77
PS	0.11	0.27	0.14	46.96	2.97
VIPS	0.07	0.06	0.06	7.42	0.91
WT	0.25	0.22	0.21	10.9	2.23

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.06	0.11	0.07	25.99	1.51
PS	0.19	0.45	0.24	46.96	5.24
VIPS	0.28	0.16	0.17	7.42	1.99
WT	0.48	0.43	0.42	10.9	4.5

So all together there are four testing variables: algorithms, datasets, HTML-type and metrics. This yielded 32 test runs in total, the results of which are presented in Table 1-4. For each algorithm we show the average Precision, Recall and F-Score values. Precision is a measure of quality that is defined as the fraction of relevant results out of all retrieved results. Recall is a measure of quantity that is defined as the fraction of retrieved results out of all relevant results. The F-Score is a combination of the two, defined as $F = 2 * \frac{P * R}{P + R}$. Additionally we also show the average number of retrieved blocks, valid hits (i.e. the number of relevant results returned by the algorithm) and the total number of relevant results (determined by the ground truth). The latter is interesting as it shows the difference in the average number of retrieved blocks and it also shows differences between the two datasets.

The Random Dataset

We first present the results of running the four different algorithms on the dataset consisting of 82 randomly selected pages. On average we have 12.24 relevant blocks on a random page. This number is the same for all the experiments with the random dataset. BlockFusion returns on average about twice as many blocks as there are relevant blocks. PageSegmenter returns about four times as many blocks as there are relevant blocks. VIPS returns too few blocks on average. Finally, WebTerrain is the closest in the number of retrieved results to relevant results. As we expected, results are considerably better for the fuzzy match metric as compared to the exact match.

The first two tables (see Table 1) show the results for the exact match metric (on the left), and the fuzzy match metric (on the right) for the random dataset with HTML input. Under the exact match metric, Precision and Recall are generally very low. BlockFusion (BF) and VIPS recognize hardly anything. Precision is highest for WebTerrain (WT) and Recall is highest for PageSegmenter (PS). Precision and Recall are clearly better for the fuzzy match metric with the number of hits roughly doubling. Especially VIPS improves substantially, indicating that a number of its blocks were only slightly off from the ground truth. The best F-Score (0.42, WebTerrain) is still rather low.

Table 2. Results for Random-DOM-Exact (left) and Random-DOM-Fuzzy (right)

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.08	0.14	0.09	30.96	1.79
PS	0.11	0.39	0.16	65.04	4.47
VIPS	0.36	0.21	0.24	9.24	2.73
WT	0.34	0.29	0.29	10.58	3.04

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.1	0.17	0.12	30.96	2.35
PS	0.15	0.51	0.2	65.04	6.12
VIPS	0.51	0.26	0.3	9.24	3.33
WT	0.57	0.49	0.49	10.58	5.33

For the DOM HTML input (see Table 2), we again observe a notable improvement when comparing the exact to the fuzzy match metric, but not quite as dramatic as for the original HTML. The number of retrieved blocks is generally higher (slightly lower for WebTerrain), reflecting the observation that the DOM HTML is typically more complex (things are added, rather than removed).

BlockFusion is performing poorly under the exact match metric, but better than on the original HTML. PageSegmenter again exhibits low precision and high recall. VIPS has the best precision but lower recall, while WebTerrain does similarly on both, giving it the best F-Score. For the fuzzy match metric, we see about a 50% improvement compared to the exact match metric. WebTerrain and VIPS have the highest precision, and PageSegmenter and WebTerrain have the highest recall. Compared to the original HTML we see some improvements as well, especially for the VIPS algorithm. Overall we see the highest scores here out of all benchmarks.

The Popular Dataset

We now turn to the the dataset consisting of 70 popular pages. On average we have about 16.1 relevant blocks on a page: 16.15 for BlockFusion, 16.22 for PageSegmenter, 16.11 for VIPS, and 16.09 for WebTerrain. The slight variation in relevant blocks is because we had to exclude a few (no more than four) pages for some of the algorithms, as they would not be handled properly due to issues in their implementation (e.g. a GTK window would simply keep hanging).

Between the original HTML and the DOM HTML one can see that the number of retrieved blocks universally goes up, giving another sign that the DOM HTML generally contains more content. Overall the results are again better for the DOM HTML, questioning the use of the original HTML in web page segmentation algorithms.

Again, we start by looking at the variations that work on the original HTML input, see Table 3. The pattern seen in the random dataset repeats: results for the fuzzy match metric are about twice as good as for the exact match metric. Both BlockFusion and PageSegmenter return decidedly too many blocks on average, but only PageSegmenter can translate this into high recall scores. VIPS and WebTerrain are fairly close to the relevant number of blocks. For the exact match metric, the results are generally poor with WebTerrain having the best precision and PageSegmenter having the best recall. The results are better

Table 3. Results for Popular-HTML-Exact (left) and Popular-HTML-Fuzzy (right)

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.03	0.06	0.03	72.85	1.07
PS	0.05	0.24	0.08	124.43	4.05
VIPS	0.07	0.09	0.07	16.72	1.17
WT	0.18	0.17	0.16	13.86	2.19

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.05	0.12	0.06	72.85	2.07
PS	0.09	0.42	0.13	124.43	6.74
VIPS	0.13	0.15	0.12	16.72	2.23
WT	0.37	0.35	0.33	13.86	4.81

Table 4. Results for Popular-DOM-Exact (left) and Popular-DOM-Fuzzy (right)

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.03	0.08	0.04	81.75	1.34
PS	0.05	0.27	0.07	163.71	4.56
VIPS	0.13	0.14	0.12	19.51	2.25
WT	0.27	0.28	0.26	14.75	3.77

Alg.	Prec.	Rec.	F-Sc.	Retr.	Hits
BF	0.04	0.12	0.06	81.75	2.13
PS	0.07	0.41	0.11	164	6.68
VIPS	0.19	0.21	0.17	19.51	3.29
WT	0.47	0.46	0.42	14.74	6.43

for the fuzzy match metric, but overall still not convincing. Again WebTerrain and PageSegmenter are the best for precision and recall respectively.

Finally, we consider the two tables for DOM HTML input for the popular dataset (Table 4). Similar to what we saw in the random dataset the improvement from exact to fuzzy matches is smaller than it was for the original HTML, but still substantial.

For the exact match metric, the results are overall better than for original HTML with the biggest gains for VIPS and WebTerrain. WebTerrain has both the highest precision and the highest recall in this test. The results for the popular dataset are the best again, as in the random dataset, when running on the DOM HTML and using the fuzzy match metric. The results for BlockFusion are again the worst. PageSegmenter has again low precision and high recall. Noticeably different is VIPS, as it does not exhibit a high precision, as it did for the random dataset. Recall is similar, though slightly lower. WebTerrain exhibits the highest precision and recall, but precision is 0.1 points lower and recall 0.03 points lower than for the random dataset.

5 Reflection

In this section, we discuss the results of the previous section in some detail. We first consider what the observed effects were of the various testing variables.

Random vs. Popular Datasets. As can be seen from our results all algorithms perform virtually always better on the random pages than on the popular pages. We believe this is due to the increased complexity of popular pages, which can be seen from the fact that they on average had 32% more blocks than a random page. Furthermore we also found that a popular page on average consists of 196.2 files in total (this number includes all the external resources referenced from a page), while a random page only consists of 79.4 files on average. The number of retrieved blocks are also universally higher for all algorithms on the popular

pages. But while the number of blocks in the ground truth was only 32% higher, the numbers for the algorithms increased by (much) more than that: BlockFusion 164.1%, PageSegmenter 152.1%, VIPS 111.1%, WebTerrain 39.3%. It thus seems that the algorithms do not scale well with increasing complexity. This could also partly explain why our results are generally less favorable than what has been found in earlier publications, as they are up to 10 years old, and the Web has become much more complex since then. It also shows the need for new techniques that deal well with this increased complexity.

Exact vs. Fuzzy Match Metric. We found that the number of recognized blocks improved significantly when using the fuzzy match metric as opposed to the exact match metric, as was to be expected. We believe that the results from the fuzzy match metric are generally more valuable since the quality of blocks will still be sufficient for most applications. Furthermore it can easily be adjusted to find more or less precise matches by adjusting the matching ratio.

Original HTML vs. DOM HTML. Comparing the original and the DOM HTML we found that the results of the segmentation for the DOM HTML are virtually always better, which is true for all algorithms on both datasets. This is due to the fact that the DOM HTML is what the user ultimately sees in the browser, it is thus the final result of the rendering process. While in the past it might have been sufficient to analyze only the original HTML, this is not true any more. As the Web becomes more dynamic and the use of Javascript to manipulate the page becomes more prevalent, there is not necessarily a link between original and DOM HTML any more. This also implies that one cited advantage of text-based segmentation algorithms, namely that they do not require the DOM to be built and are thus very fast, is not true any longer, as even for these algorithms it is necessary to obtain the final HTML for optimal results.

The Four Segmentation Algorithms. The four algorithms differ widely in their performance, and none of them performed well enough to be universally applicable, as the highest average F-Score was 0.49 (WebTerrain). Our comments here pertain to the test runs using the fuzzy match metric and the DOM HTML because we consider those the most relevant. But the general conclusions hold for the other testing combinations as well.

BlockFusion. This algorithm showed the worst performance on both datasets. Both precision and recall are very low (< 0.1 and < 0.2 respectively). It also returns too many blocks on average (2.5x too many for the random dataset and 5.1x too many for the popular dataset). We could thus not repeat the results from [13]. We conclude that a solely text-based metric is not sufficient for a good segmentation, but that it can be used to augment other approaches.

PageSegmenter. This algorithm exhibits low precision and (relatively) high recall (< 0.2 and > 0.4 respectively). This is due to the fact that it retrieves by far the most blocks of all algorithms (5.3x too many for the random dataset and 10.1x too many for the popular dataset). The number of false positives is

Table 5. ANOVA summary

	Df	Sum Sq	Mean Sq	F value	p-value
algorithm	3	3.55	1.18	333.41	; 2e-16
html	1	0.35	0.35	98.24	; 2e-16
dataset	1	0.37	0.37	104.55	; 2e-16
metric	1	0.66	0.66	185.48	; 2e-16
algorithm:html	3	0.22	0.07	20.97	2.10e-13
algorithm:dataset	3	0.07	0.02	6.59	0.000198
algorithm:metric	3	0.12	0.04	11.66	1.38e-07
html:dataset	1	0.05	0.05	14.61	0.000136
html:metric	1	0.02	0.02	5.60	0.018055
dataset:metric	1	0.00	0.00	0.55	0.456698
Residuals	2298	8.15	0.00		

thus very high. It would thus be interesting to see if this algorithm could be optimized to return fewer blocks while retaining the good recall rates.

VIPS. This algorithm showed the biggest difference between the random and the popular dataset. Precision was high and recall mediocre for the random dataset (0.51 and 0.26 respectively), while both were low for the popular dataset (0.19 and 0.21 respectively). It is not clear why there is such a substantial difference. The number of retrieved results is slightly too low for the random dataset, while it is slightly too high for the popular dataset (25% too low and 21% too high respectively). In terms of the F-Score the VIPS algorithm was second-best.

WebTerrain. This algorithm showed relatively high precision and recall for both datasets (both > 0.4). It retrieved slightly too few blocks for both datasets (14% too few for the random dataset and 8% too few for the popular dataset). We thus find that a combination of structural and rendering-based approaches enhances overall results. Furthermore the terrain heuristic seems promising. Future work could therefore likely improve upon these results by using more sophisticated combinations of different approaches and heuristics.

Analysis of Variance. We performed an analysis of variance (ANOVA) to test the impact of the four factors `algorithm` (with levels: `blockfusion`, `pagesegmenter`, `VIPS`, `WebTerrain`), `html` (levels: `dom`, `html`), `dataset` (levels: `popular`, `random`) and `metric` (levels: `exact`, `fuzzy`) on the F-score; see Table 5.

To reduce the non-normality of the dependent variable, we performed a Box-Cox transformation of the F score. The transformation is defined as:

$$Y^{(\lambda)} = \begin{cases} \frac{Y^\lambda - 1}{\lambda} & \text{for } \lambda \neq 0 \\ \ln Y & \text{for } \lambda = 0 \end{cases}$$

where we defined $Y = Fscore + 1$ because Y is required to be strictly positive. We selected the value of λ that yielded the highest log-likelihood value, using a simple grid search [9]. This resulted in $\lambda = -3.45$.

We can read from the column labeled *p-value* that for all terms except `dataset:metric` the null hypothesis of 'no effect' will be rejected at the conventional significance level of $\alpha = 0.05$. To avoid confusion we note that the Analysis of Variance was performed on the (transformed) *F scores* obtained for different combinations of factor levels, which is an entirely different quantity than the *F statistic* (reported in the column *F value* in Table 5) which is used to test the significance of different (combinations of) factors in explaining the variation in observed *F scores*.

We see in Table 5 that the four factors `algorithm`, `html`, `dataset` and `metric` are all highly significant ($p\text{-value} < 2e\text{-}16$) in explaining observed variation in F-scores. The analysis thus confirms our intuition that these factors are relevant for an analysis of web page segmentation algorithms.

We also included *interaction terms* in the analysis (all the colon-separated variables, such as `algorithm:html`). An interaction term is the product of two variables that can in itself be a significant predictor of the outcome. A high significance for an interaction term $x : y$ means that x and y interact, i.e. the effect of x on the outcome (i.e., the F-score) is different for different values of y .

In Table 5 we see that all interaction terms are significant at $\alpha = 0.05$, except for the term `dataset:metric`. The term `dataset:metric` not being significant means that the influence of `metric` on the F-score is typically similar for both popular and randomly selected pages. This is consistent with the results reported in Section 4, where we found that the fuzzy metric always returns a higher F-score than the exact metric, for both randomly selected and popular web pages.

Threats to Validity. As in any empirical study there are various threats to validity. Unless noted otherwise, all discussed threats are to external validity.

A threat to construct validity concerns the implementations and our interpretations of the algorithms. Because we had to make these interpretations, and we did not obtain answers to our e-mails about these interpretations from the authors, we run the risk that we are not exactly measuring their algorithm, but a variation thereof. Detailed discussions of these issues can be found in [15].

In our study, we compare approaches to web page segmentation by looking at a particular instance of each approach. In our selection we have chosen well-known, often-cited representatives of each approach. To compensate, we have made our framework and datasets open for everyone to use, so that others can easily extend upon our work, by implementing other instances of the paradigms testing these against the ones that we have implemented.

When it comes to the experimental data, we used two datasets: one with popular web pages, mainly because doing well on pages that are often read by people is something an algorithm should be rewarded for, but also to serve as a "worst-case" since we expected these websites to be more complex than the average website. Since we also did not want to bias too much towards such pages, we also included a large sample of random pages, with the aim of increasing external validity. However, we did have to drop 48 out of 200 web pages, because when rendered locally they differed from the original web page. This means that

our results may not generalize to websites that modify the DOM extensively, using Javascript.

The mark up was performed according to a particular ontology (Section 3), which may hurt external validity. We do believe our ontology to be generic enough to be applicable to most existing websites. A second issue at this point is construct validity: did the participants who marked up the websites for us understand the ontology, and what was expected of them? To make sure that that was the case, the first author first explained the ontology to them, and checked five segmentation results for each participant to see whether they had understood him well enough. It was also verified for a few sample pages that the volunteers agreed on the web page segmentation for those pages. We note that our test subjects all have an IT background, which may decrease external validity.

An issue in our study is that one of the algorithms we consider has been of our own devising. We note, however, that having confirmed that our implementation was correct, i.e., it behaved as we designed it to do on a few web pages, we did not make any modifications to it during or after we ran our experiments. This to avoid the danger of overfitting our algorithm to the chosen datasets.

6 Related Work

A comprehensive discussion of related work can be found in [15]. Here, we restrict ourselves to a selection that is most closely related to our work here.

The research on structuring information on web pages into semantic units goes back to at least [11]. Subsequent authors tried to automate the process of locating where the information of interest resided. The process can be broken into two distinct steps: segmentation (what belongs together in blocks) and labeling (what is the best description for the block, in terms of some chosen ontology).

In [14], the approach is based on heuristics that take visual information into account. They built their own basic browser engine to accomplish this, but do not take style sheets into account, and they avoid calculating rendering information for every node in the HTML tree. They then define a number of heuristics on the rendered tree assuming that the areas of interest on a page are header, footer, left menu, right menu and center of the page, and where they should be located, e.g. header on top. The authors test their algorithm experimentally for a dataset for which they manually label areas on 515 different pages. Their overall accuracy in recognizing targeted areas is 73%.

In [7], the authors turn the DOM tree into a complete graph, where every DOM node is a vertex in the graph. Each edge is then assigned a weight that denotes the cost of putting these two vertices into the same segment. The weights are learned from a dataset regarded as the ground truth by looking at predefined visual- and context-based features. Finally they group the nodes into segments by using either a correlation clustering algorithm, or an algorithm based on energy-minimizing cuts; the latter performs considerably better. Their evaluation is based on manually labeled data (1088 segments on 105 different web pages).

Baluja [2] focuses on the application of optimizing existing web pages for mobile phones by, first, dividing the web page into a 3x3-grid. The user can then interactively arrange for the website to be optimized for mobile phone screen. The page segmentation algorithm is based on clues from the DOM combined with a number of computer vision algorithms. Specifically, they use an entropy measurement to construct a decision tree that determines how to segment the page. They test their approach on a number of popular websites where they achieve good results in most cases (they rarely cut through coherent texts). One artificial limitation of their approach is that it divides the page into at most 9 segments, but it seems possible to adapt it to other grid sizes.

In [1], the authors improve upon the VIPS algorithm, by improving the first phase of visual block extraction. They divide HTML tags (including the ones introduced by HTML 5) not into three classes but into nine instead, and define new separation rules for these classes based on visual cues and tag properties of the nodes. No empirical evaluation of their algorithm is provided.

We found only one paper that, like us, is focused on comparing existing approaches to web page segmentation and labeling: [18]. The classification of about 80 papers is largely qualitative including bottom-up vs. top-down, DOM-based vs. visual, how the evaluation of the approaches is measured (precision and recall, success rate, or execution time), and whether specific heuristics are employed based on assumptions about the layout of web pages.

7 Conclusion and Future Work

We have compared the performance of four algorithms for web page segmentation. Our work shows that the three older algorithms BlockFusion, PageSegmenter and VIPS, performed worse than they did in their original publications, something which we believe is due to increasing complexity of websites and their ever more dynamic behavior due to the growing prevalence of DOM manipulations via Javascript. This belief is affirmed by our finding that the algorithms do better on randomly selected websites than on the more complicated popular ones. It does seem that combining two approaches to segmentation is a promising direction for further research, as shown by the consistently highest F-scores of the WebTerrain algorithm, although more research is needed to make (automatic) web page segmentation practical. Promising-looking directions are more sophisticated combinations of different approaches and more directed segmentation algorithms that e.g. only focus on certain segments on a page or that target only specific domains of websites.

References

1. Akpınar, E., Yesilada, Y.: Vision based page segmentation: extended and improved algorithm. Technical report, Middle East Technical University Northern Cyprus Campus, January 2012

2. Baluja, S.: Browsing on small screens: recasting web-page segmentation into an efficient machine learning framework. In: Proceedings of the 15th International Conference on World Wide Web (WWW 2012), pp. 33–42. ACM Press (2006)
3. Boldi, P., Rosa, M., Santini, M., Vigna, S.: Layered label propagation: a multiresolution coordinate-free ordering for compressing social networks. In: Proceedings of the 20th International Conference on World Wide Web (WWW 2011). ACM Press (2011)
4. Boldi, P., Vigna, S.: The WebGraph framework I: compression techniques. In: Proceedings of the Thirteenth International World Wide Web Conference (WWW 2004), pp. 595–601. ACM Press, Manhattan (2004)
5. Cai, D., Yu, S., Wen, J.-R., Ma, W.-Y.: Extracting content structure for web pages based on visual representation. In: Zhou, X., Zhang, Y., Orlowska, M.E. (eds.) APWeb 2003. LNCS, vol. 2642, pp. 406–417. Springer, Heidelberg (2003)
6. Cai, D., Yu, S., Wen, J.R., Ma, W.Y.: VIPs: a visionbased page segmentation algorithm. Technical report, Microsoft Technical Report, MSR-TR-2003-79 (2003)
7. Chakrabarti, D., Kumar, R., Punera, K.: A graph-theoretic approach to webpage segmentation. In: Proceedings of the 17th International Conference on World Wide Web (WWW 2008), pp. 377–386. ACM Press (2008)
8. Crawled by the Laboratory of Web Algorithmics, University of Milan, <http://law.di.unimi.it/>. Yahoo! research: “Web spam collections”. <http://law.di.unimi.it/webdata/uk-2007-05/>
9. Fox, J.: Applied Regression Analysis and Generalized Linear Models. Sage, 2nd edition (2008)
10. Guha, R., McCool, R.: TAP: a semantic web test-bed. *Web Semantics: Science, Services and Agents on the World Wide Web* 1(1), 81–87 (2003)
11. Hammer, J., Garcia-Molina, H., Cho, J., Aranha, R., Crespo, A.: Extracting semistructured information from the web. In: Proceedings of the Workshop on Management of Semistructured Data, pp. 18–25 (1997)
12. Kohlschütter, C., Fankhauser, P., Nejdl, W.: Boilerplate detection using shallow text features. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM 2010, pp. 441–450. ACM, New York (2010)
13. Kohlschütter, C., Nejdl, W.: A densitometric approach to web page segmentation. In: Proceedings of the 17th ACM Conference on Information and Knowledge Management, pp. 1173–1182 (2008)
14. Kovacevic, M., Diligenti, M., Gori, M., Milutinovic, V.: Recognition of common areas in a web page using visual information: a possible application in a page classification. In: IEEE International Conference on Data Mining (ICDM 2002), pp. 250–257 (2002)
15. Kreuzer, R.: A quantitative comparison of semantic web page segmentation algorithms (MSc thesis) (2013). <http://www.cs.uu.nl/wiki/Hage/SupervisedMScTheses>
16. Vadrevu, S., Gelgi, F., Davulcu, H.: Semantic partitioning of web pages. In: Ngu, A.H.H., Kitsuregawa, M., Neuhold, E.J., Chung, J.-Y., Sheng, Q.Z. (eds.) WISE 2005. LNCS, vol. 3806, pp. 107–118. Springer, Heidelberg (2005)
17. Vadrevu, S., Velipasaoglu, E.: Identifying primary content from web pages and its application to web search ranking. In: Proceedings of the 20th International Conference on World Wide Web (WWW 2011), Hyderabad, India (Companion Volume), pp. 135–136. ACM Press (2011)
18. Yesilada, Y.: Web page segmentation: A review. Technical report, Middle East Technical University Northern Cyprus Campus, March 2011