Metrics for the Evaluation of News Site Content Layout in Large-Screen Contexts

Michael Nebeling
Institute of Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland
[lastname]@inf.ethz.ch

Fabrice Matulic
Institute of Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland
[lastname]@inf.ethz.ch

Moira C. Norrie
Institute of Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland
[lastname]@inf.ethz.ch

ABSTRACT

Despite the fact that screen sizes and average screen resolutions have dramatically increased over the past few years, little attention has been paid to the design of web sites for large, high-resolution displays that are now becoming increasingly used both in enterprise and consumer spaces. We present a study of how the visual area of the browser window is currently utilised by news web sites at different widescreen resolutions. The analysis includes measurements of space taken up by the article content, embedded ads and the remaining components as they appear in the viewport of the web browser. The results show that the spatial distribution of page elements does not scale well with larger viewing sizes, which leads to an increasing amount of unused screen real estate and unnecessary scrolling. We derive a number of device-sensitive metrics to measure the quality of web page layout in different viewing contexts, which can guide the design of flexible layout templates that scale effectively on large screens.

Author Keywords

Web Design Guidelines, Web Site Adaptation to Large Screens, Technical Tools for Usability Evaluation

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Screen design*

General Terms

Design, Human Factors

INTRODUCTION

The popularity of hand-held devices with web browsing functionality has prompted media publishers to design special, accessible web sites for these devices. Online news sites in particular have striven to address customer demand in delivering content that can be consumed on smartphones and the like. However, at the other end of the spectrum, little is being done to cater for the range of large-display surfaces that are now in widespread use in offices, homes and public spaces.

As shown by recent reports from the market research firm DisplaySearch[1], average screen dimensions and resolutions are increasing; moreover, the trend is towards displays with a 16:9 format, especially for notebook computers. In recent years, there has also been a growing number of other types of large-display devices such as wall screens, digital tabletops and projected surfaces, all predominantly geared towards horizontal, wide-format aspect ratios. This change in the browsing device landscape requires adaptations of many web site layouts where the often strongly vertical orientation at fixed width seems a poor design choice for such device characteristics. Also, with higher resolutions, the text becomes physically smaller and therefore less readable. Some web sites attempt to alleviate this problem by including a function allowing users to change the size of the font, an effect that can also be achieved using zoom features available in modern browsers. Another common function is a printer-friendly view where only the main content is shown with minimal or no layout constraints so that text can flow freely to fill the available space. However, those are only minor improvements of the user experience and the fact that they require user intervention is not always convenient.

We therefore decided to investigate both the technical and design issues raised by adaptation to large-display devices, which poses new challenges such as the efficient 'upscaling' of content, not to waste valuable screen real estate and impose unnecessary scrolling on the user. We argue that it is necessary to provide flexible screen layouts that make appropriate use of the higher amount of screen space available on large displays and, as a first step, we have developed a set of metrics that can be used to assess the presentation of content in different widescreen settings and hence to inform web page design also for large, high-resolution screens. The proposed metrics are grounded in a comprehensive study of existing news web sites to see how the browser viewport is currently utilised by different news content elements at viewing sizes larger than 1024x768 and to identify key issues. In this paper, we present the results of this study and the proposed metrics as well as how these can be implemented based on

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

[1] http://www.displaysearch.com

native web technologies to provide developers with the technical tools for evaluating web page layouts in terms of the spatial coverage of content in different widescreen contexts.

RELATED WORK

Many different usability evaluation methods have been specifically crafted for the web [11], e.g. design inspections to examine the hypertext specification, web usage analysis to understand user behaviour, and heuristic evaluation of prototypes at different design stages as well as of the final web site [18]. Given today's proliferation of web-enabled devices, the usability of a web site is also determined by its ability to adapt to the specific device in use. Much attention has been paid to the adaptation of web content for mobile phones and other kinds of small-screen devices. A number of automatic methods for partitioning web pages into smaller, semantically related units that fit into the screen of a mobile device have been proposed, e.g. [3, 10]. On the other hand, for large screens that are now increasingly used in enterprise and also consumer spaces, techniques to aggegrate content from multiple related pages into a single view might be welcome to reduce unnecessary navigational overhead. However, optimal design for large-display viewing contexts remains largely unexplored. Research in HCI currently seems to focus on new interaction techniques on such devices where touch and gesture-based modalities seem to be predominant, but efforts to design user interfaces that make appropriate use of the greater amount of screen real estate are minimal. We believe that, as a first step, some sort of device-sensitive web design metrics are required that can inform current web design processes so that they also cater for the characteristics of large-display devices.

Early research proposed several metrics to quantify usability factors such as the total word count in a page, the number of links and media as well as the spectrum of colours and font styles employed [13]. It was also found that web design is a moving target and that many of the established guidelines change over time [12], which may explain the lack of concrete thresholds for a wider class of quantitative web page measures. Existing guidelines such as WCAG, the Web Content Accessibility Guidelines[2] by W3C, consist of a set of recommendations on making content accessible, primarily for disabled users, with only some advice on designing for highly limited devices such as mobile phones. Their web site also includes a list of many different evaluation and issue reporting tools, but none of them address design deficiencies related to the spatial coverage of content with respect to different viewing contexts.

While several productivity benefits have already been characterised for large-display environments, e.g. in [4, 23], there is still a lack of concepts and tools to take advantage of large screens [20]. Due to the many fixed designs with often a vertical flow, using web sites in widescreen contexts generally does not bring any real advantages to users, which is why large screens at this stage still seem to be primarily suited for multi-task scenarios. For a single application to capitalise on the greater amount of screen space, various forms of adaptation would be required that are currently not supported by most web sites. Research has shown that selecting and placing different versions of text and media to display only the relevant parts on small screens or a much higher level of detail on large displays is generally possible [1, 14, 21], but requires substantial technical overhead, especially if one was to rely on native web technologies. One can hope that a flexible segmentation of content will become easier and more popular if W3C's proposals to include new tags in HTML5[3] for the annotation of page elements as well as multi-column layout in CSS3[4] are adopted and implemented in all major web browsers. However, our ongoing experiments with these new technologies show that even the latest additions to HTML and CSS are insufficient for effective adaptation of content and presentation using only features of those web standards.

One of the key challenges for adaptive layouts is to reduce the amount of scrolling required by fitting most of the content on the screen area above the fold (i.e. the screen space that is visible at first without scrolling), while, at the same time, not overloading the page with too much information. Directly related to this aspect is the physical layout of text which is found to have a considerable impact on the on-screen reading experience [5]. There is an increasing number of relevant tools, such as Instapaper[5] or Apple's Safari Reader technology, that extract long text paragraphs from web pages and show these using font settings and a display format optimised for reading, but also they provide poor results in large-screen environments due to the common one-column, fixed vertical layout that only makes use of a minimal portion of the screen in wide-format settings.

To generally improve users' reading comfort and efficiency, many findings suggest the use of multiple text columns and some even argue for horizontal scrolling, e.g. [8, 2]; however, this brings up two important issues. The first problem is that current support for multi-column layout in popular browsers is at best experimental and subject to change, as this is a feature proposed to be included only with the upcoming CSS3 modules. The current workaround among web developers is to specify adaptive layouts with the help of JavaScript libraries such as jQuery Masonry[6], which can allow the content to be spread in multiple columns, the number and size of which depends on the dimensions of the target viewing size. However, while such flexible layout engines based on automatic methods are also desirable for printer-friendly versions of news articles to improve on the simple sequential positioning of content used by most current browsers, they give less control over the concrete design in a particular viewing context and customisation of the layout mechanisms generally requires considerable programming skills. The second issue is that vertical scrolling is predominant on most web pages and further compounded by the fact that there is often only a vertical mouse wheel, so a com-

[2]http://www.w3.org/TR/WCAG20/

[3]http://www.w3.org/TR/html5/

[4]http://www.w3.org/TR/css3-multicol/

[5]http://instapaper.com

[6]http://www.smashingmagazine.com/2009/06/09/smart-fixes-for-fluid-layouts/

plete change to horizontal scrolling would be rather radical and require many users to change their browsing habits.

In view of these problems, we have started to investigate the technical and design challenges for developing adaptive web interfaces that also cater for large-display, high-resolution devices while building on established standards. Our first results comprise the metrics presented in this paper that we have derived from an empirical study of current news site designs and how these adapt in different widescreen settings. The proposed metrics mark the first step of our effort to build a research framework that seeks to establish design guidelines as well as to form sound technical tools for the development of flexible web interfaces that also adapt to novel forms of large displays, i.e. not only larger desktop screens, but also digital TVs, tabletop systems or wall-size displays.

The study presented in the next section provided important input for developing the metrics following this. Many other studies have been conducted to understand users' browsing habits and identify patterns of navigation on the web, e.g. [22]. Several studies have also investigated the visual aesthetics of web sites as perceived by users [15] and how framing effects may influence users' web site quality judgement [9]. However, the effects of different device characteristics in terms of screen size and resolution and their potential impact on the browsing experience have so far played only a minor role in studies and have primarily been investigated for small-screen devices, e.g. [6]. For large screens, there is still only limited insight obtained from a few formative evaluations of user expectations when browsing web pages on high-resolution, wide-format displays [7]. The results there also suggest the use of constraint-based liquid layout for web pages to automatically fill larger portions of the screen as more space becomes available. However, an important factor that has not received enough attention is a definite strategy to specify suitable constraints for those adaptive layouts as these are key for such layout mechanisms to be effective. In that sense, we believe that our work is complementary and makes an important step forward, as the metrics and visual tools presented in this paper are devised to inform web page design and help developers find device-related layout constraints.

STUDY

Our study aimed at providing numbers on the status quo of text-centric web content layout and window space occupation in widescreen environments. In that regard, the closest to our study is [16] that presents a detailed breakdown of the different web page elements to relate the relative page coverage taking into account text, images and other embedded objects. However, their content analysis does not address spatial distribution of content elements at different viewing sizes, nor does it resolve the different semantics of content, for example, to differentiate actual content from ads.

The choice of news articles over other types of web sites was motivated by their relative homogeneity in terms of structure and appearance and the fact that online newspapers and magazines can be easily compared to their print counterparts. Moreover, news web sites are backed by companies that have a vital interest in delivering an optimal reading experience to their customers in order to keep up their businesses.

Methodology

In order to assess the spatial distribution of web content for a news page, we approached the problem by first identifying the major functional categories of structural elements that typically appear on a page showing a full news article. However, since the purpose of the study was to measure space occupation rather than establishing the full portrait of a web page, we decided to limit ourselves to four categories: background, article content, ad and web interface, the latter including links to other news items related or not related to the current article. The idea was then to calculate the positions and combined areas of these elements in several viewing contexts and based on those results compute a number of indicators to present the overall situation.

The web pages were all viewed and analysed on Firefox using standard settings and a set of extensions to perform the measurements. To obtain meaningful results reflecting a sufficiently representative collection of news articles, the 50 most popular news web sites as ranked by Alexa were considered[7], excluding news aggregators that only show short text snippets for each news item. For each site, a sample article was chosen and using a customised version of the Aardvark extension[8], the DOM elements of the page were marked manually according to their semantic function. Typically, the main structural blocks in a news web page have their own HTML container elements with identifying id or class attributes that are consistent across the whole web site. Those attributes were used to associate marked areas with their semantic function so that other articles retrieved from the same web site could then be broken down and tagged automatically. All pages that were marked in this way were subsequently viewed by the authors to check for inconsistencies and discrepancies that would distort the measurements. Some web sites split long articles over several pages and, in those cases, values were added or averaged as appropriate to match the criteria used for single-page layouts. Any reader comments were also hidden or removed as they do not belong to the article per se, although technically an integral part of the page.

The measurements consisted of determining the position and length of marked blocks and calculating the combined area that they covered in the page. Each article was loaded simultaneously in five independent browser windows set to different sizes. The dimensions were chosen based on browser resolution statistics provided by W3Counter[9]. The three most popular widescreen resolutions were selected along with the reference 1024x768 resolution for 4:3. The resolution 1680x 1050 was also reversed for one measurement set so as to give the window a vertical orientation in order to simulate a rotating monitor in portrait mode. Since the study focused on raw, pixel-based space occupation, the physical screen

[7]http://www.alexa.com/topsites/category/Top/News

[8]http://karmatics.com/aardvark

[9]http://www.w3counter.com/globalstats.php

size did not come into consideration at that point. The latter is however important when taking into account subjective, user-dependent criteria such as reading comfort and needs to be addressed when designing web pages for large screens.

Each browser window was left in its default appearance with only the navigation and bookmark toolbars. This means that the actual dimensions of the browser viewport were $w - 6px \times h - 136px$, including possible scrollbars. The sizes reflect a situation where the browser window is maximised or set to full screen, which according to many user studies (e.g. [22]) is what the majority of users adopt. For situations where the browser is run in windowed mode and hence at intermediate user-defined window sizes, it should be easy to extrapolate from the results for the considered resolutions.

Results

The vast majority of news web sites currently adopt a two or three-column layout with one or more headers and footers, and the article content itself in the left or central column. Navigation items such as menus, buttons and links to related content typically appear in sidebars surrounding the article, at the top and directly below it in the footer(s). Frequently, floating blocks also appear inside the main article causing breaks in the article's text flow and formatting. As for the ads, while their content frequently changes between loads, they are generally placed at fixed locations and have fixed dimensions for a given web site. No particular design rule seems prevalent as to where the ads should appear. Ad banners and inserts can be embedded in virtually any of the aforementioned elements, including within the article text itself. With only a few exceptions, the web page content is centred, except perhaps for some footers, headers or background elements that are made to stretch to the whole width of the window. The area covered by the article content itself in the selected pages averaged 42% of the total web area and the space taken up by ads roughly 12%. This perhaps contrasts with blogs and other user content-intensive sites, where entries are typically appended one after the other inside a long column (and where possibly there is even more space lost due to sub-optimal layout strategies).

Figure 1 shows the average percentages of the areas covered by web, article content and ad elements in the client space, i.e. the total area of the browser window, including visible and invisible parts. As the resolution increases, the amount of space used by actual web content decreases, except when the window is set to vertical mode (1050x1680). A side effect of this phenomenon is the appearance of margins on each side that become larger as the viewport width increases, as illustrated in Fig. 2. The margin lengths are computed by taking the leftmost and rightmost pixels of the first and last solid elements of the web document (i.e. non-background) at each scanline and determining the distances to the left and right edges of the viewport, respectively. (If the document width is larger than the viewport then the margin is negative). The average margin width is then calculated for the whole page. This means that pages with long articles extending beyond their flanking sidebars would see their line margin widths significantly increase if columns are static and content is not reflowed inside them. At screen widths around 1024 pixels, the margin is minimal and there is a jump above 1200. This is due to the fact that many web sites are currently designed for a target resolution of 1024x768. But, even at this size, the margin average is more than 10% due to unbalanced column heights. As a matter of comparison, the amount of white space used by a page of a regular print newspaper is roughly 10% of the paper surface (mostly covered by margins).

Another consideration which perhaps matters to publishers and editors is what users see when the web page is first loaded in the web browser and the topmost portion of it ap-

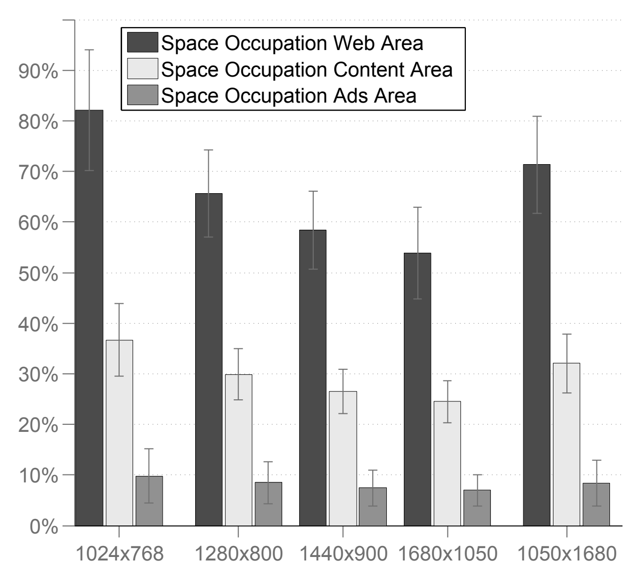

Figure 1. Area coverage in whole client space

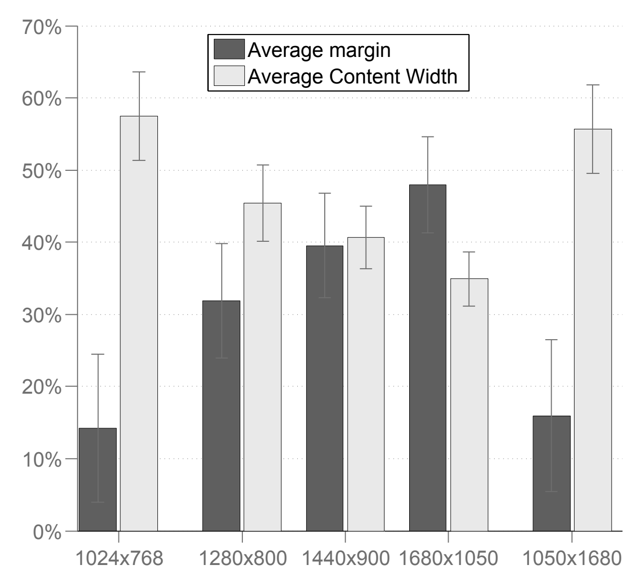

Figure 2. Average margin and content widths as percentage of the total viewport width

pears in the window. Too many extraneous elements, ads and too much background might indeed frustrate readers who are mostly interested in viewing the article. For publishers sustaining their business through advertisements, the placement of ad containers on the web page might be an important design factor. If ads are located in areas which readers never come to see, then this represents a potential loss of revenue. As it turns out, a significant proportion of ads are placed below the article (about 28% of the total ad surface), a zone that may never be exposed to users that do not scroll that far.

Figure 3 depicts the area distribution of the different categories visible on the area above the fold (or 'first screen'), where the 'nav' area is the area that remains after content and ad areas have been subtracted. Figure 4, on the other hand, shows the extent of surface coverage by the web page's elements on the first screen as a percentage of the total space they occupy in the whole web document (here 'web area' includes article content and ad areas). The figures show once again that, as widescreen resolutions increase, the browser window is mainly filled with background pixels at the detriment of content. At 1680x1050, the user only sees less than 1.5 times the amount of a page's content than at 1024x768 on average, although there is more than 2.2 times as much available space. Viewing the page at 1050x1680 yields much better results, as the user is able to see more than half of the web page and almost 60% of the article content without scrolling.

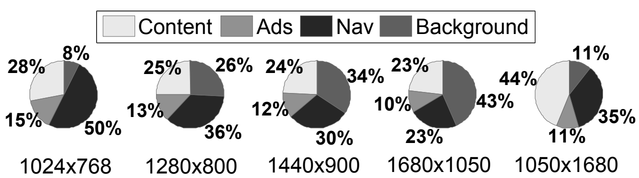

Figure 3. Area distribution on first screen client space

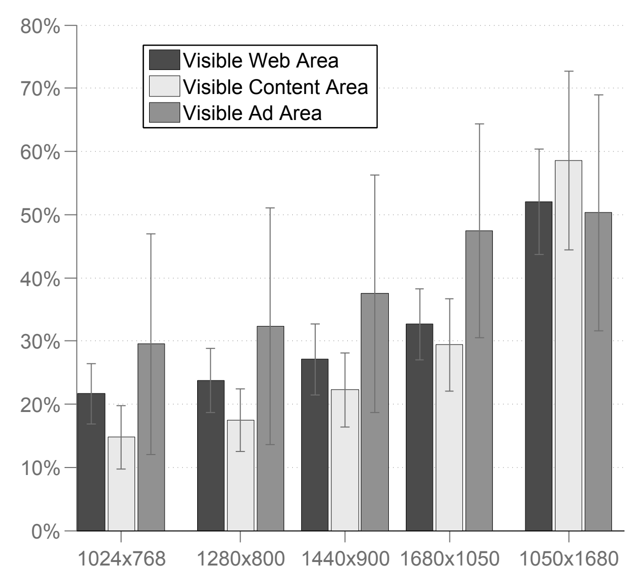

Figure 4. Percentage of visible elements on first screen from total surface in page

The same effect can be observed with ads, which generally do not adapt well to higher amounts of screen real estate. While they still occupy 15% of the first screen at 1024x768, the visible ad area decreases to 10% at 1680x1050. Even though much more content becomes visible at its counterpart resolution of 1050x1680, the visible ad area only slightly increases to 11%.

A simple metric to measure how much space could be saved if screen real estate were used with maximum efficiency is to compare the amount of scrolling needed to view the whole page as rendered with the amount required to display the document's foreground elements when optimally filling the available space. Figure 5 shows the scroll factors of the two theoretical fills (web and article content) vs. that of the actual rendering expressed as the number of viewports needed to show the total area of the considered categories. A value

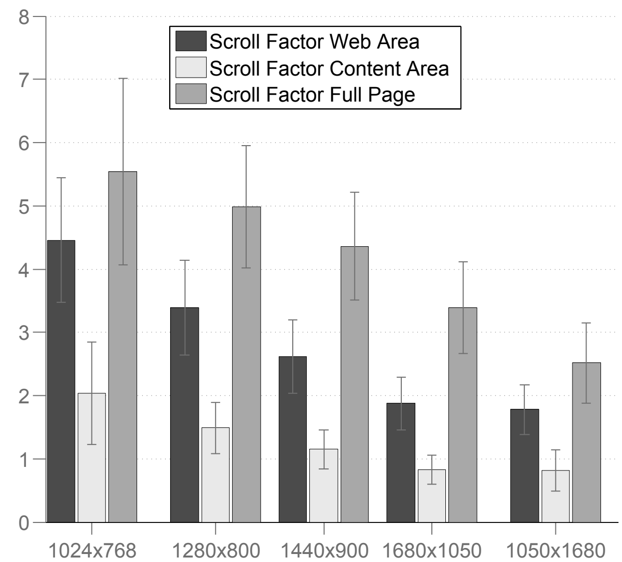

Figure 5. Scroll factor for optimal web and content coverage vs. actual rendering

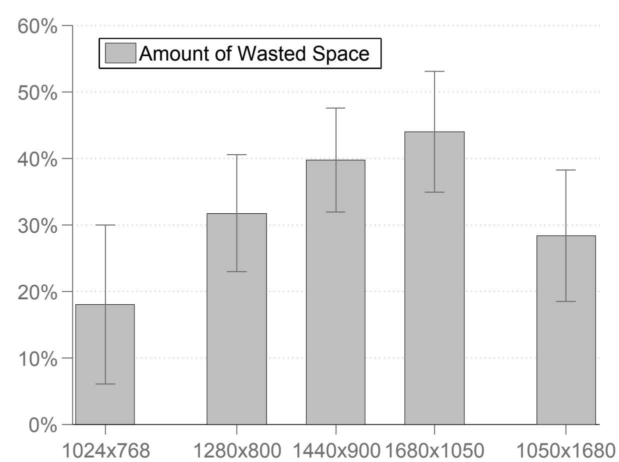

Figure 6. Amount of space that could theoretically be saved by optimal use of screen real estate

lower than 1 means that the entire area can be displayed in the viewport and no scrolling is required. The results for the content areas reflect a hypothetical case where only the actual content of the article itself is shown on the screen, similar to printer-friendly renderings.

Juxtaposing the theoretical scroll factors with the values for the actual renderings reveals increasing gaps between the optimal and real measures. At higher widescreen resolutions, more and more unnecessary scrolling is imposed on the user to view the page's entire content. The amount of space that is potentially 'wasted' is summarised in Fig. 6. Once more, the data shows a clear increase of unused screen real estate in high-resolution horizontal contexts, as web pages generally fail to adapt to the change of rendering environment.

Discussion

Obviously, web design is more than simply optimising space usage and the costs of leaving much of the screen bare on a large display are not as critical as those associated with wasting paper in a print newspaper where costs are real. One could even say that the penalty of a web page with an ill-suited layout is only the fact that it involves more scrolling and clicking for the reader. However, as the diversification of browsing clients further expands, the need for more flexible layouts to accommodate more demanding users will no doubt arise.

Discussions of the respective merits of static layouts versus flexible or liquid layouts have sparked lively debates among designers. Proponents of the former point out the ease and convenience of designing for an area with fixed dimensions, while the advocates of the latter argue that accessibility should be the prime factor in content presentation [17]. Print newspapers and magazines make near-optimal use of the available paper surface using columns and separations, but their template sizes are fixed and the appearance of the final product is determined at printing time. To achieve optimal surface occupation without sacrificing aesthetic considerations of layout design on digital surfaces, whose dimensions and resolutions are unknown at creation time, is a non-trivial engineering challenge.

One argument against flexible layouts is that they often allow text column widths to stretch without limit, thus producing excessively long lines that impose cognitive strain on the reader when trying to follow on to the next line. Among the reviewed web pages, the maximum number of characters per line for article columns with fixed widths was on average 94. This already exceeds the recommended 70 characters per line for pleasant and efficient reading [19], and increasing column widths, and hence line lengths, further would detract from readability and the overall user experience.

Another decisive factor that heavily influences readability is typography, i.e. font type and size [5]. News sites tend to adopt either modern-looking sans-serif fonts, mostly Arial, or a serif font such as Georgia or Times New Roman. The latter seem to be particularly favoured by traditional newspapers such as The New York Times and The Washington Post, perhaps to retain a distinctive print-like look to assert their identity. The typical font size of paragraphs in news articles tends to be between 13 and 15 pixels, regardless of the viewing resolution. As mentioned in the introduction, there are several ways to compensate for text that is rendered too small by default, but they typically require user intervention. While one can imagine site-specific text-display settings optimised for a particular viewing environment, it should be possible for web sites to automatically adapt the default font settings based on thresholds relative to the increase in resolution even with no user preferences set.

METRICS

The problems identified in our study essentially relate to three major aspects of web layouts: (1) the overall use of the screen, (2) the proportions between different content elements and (3) the readability of larger amounts of text. We can trace the cause of the observed problems to the following design issues.

- Many web layouts have been designed for a fixed width only. In widescreen contexts, the greater amount of screen space is therefore potentially wasted as it gets mostly filled with background content.
- Text and font characteristics such as font weight and size, number of columns and column widths as well as line spacing are often optimised for the 'standard' screen resolution of 1024x768 and typically hard-coded. As a consequence, text flow and style hardly scale with larger resolutions to maintain readability in widescreen contexts.
- Media resources such as images, videos and animations are often available only in one size and embedded with fixed dimensions. As a result, the embedded media do not scale well in widescreen contexts. For instance, ads that make extensive use of animated GIFs or Flash animations to draw the attention of the user consequently lose high proportions of the visible area on the first screen in comparison to the text.
- While web sites generally attempt to place the most important content elements on the first screen, the amount of information presented to the user at larger viewing sizes often increases substantially as more content becomes visible at once. Particularly for the navigation options, this can lead to an overflow of visible links, which makes it harder for users to find a particular navigation point.

In order to quantify these issues, we have defined the metrics shown in Table 1. The basis for these metrics build the window area, web area and content area. Similar to our study, the window area is the browser viewport and therefore the room available for the content elements to be visible without scrolling. The web area is the space that all foreground and background content elements in the document require to be viewed at the same time. The content area refers to the space occupied only by the foreground content elements. Depending on individual web site requirements, it can be helpful to divide the content area into groups of semantically related content elements in order to compare how often elements of

Table 1. Metrics to measure the quality of web layouts with respect to the viewing context

Metric	Description
Document-window ratio = $\frac{\textit{Web area}}{\textit{Window area}}$	Determines how many viewports are required to fit the whole document. This ratio can be divided into vertical and horizontal scroll factors determined by height and width ratios, respectively. Web sites at larger viewing sizes should approach a value closer to 1. A value less than or equal to 1 means that no horizontal and/or vertical scrolling is required, but very low values mean that most of the available space is not used.
Content-window ratio = $\frac{\textit{Content area}}{\textit{Window area}}$	Determines how many viewports are required to fit the foreground or actual content only. Significantly different values between the document-window and content-window ratios for larger viewing environments hint at potentially 'wasted' screen real estate, as the space that becomes available would primarily be consumed by background content.
Wide text ratio = $\frac{\textit{Wide text area}}{\textit{Total text area}}$	Determines how much of the total text area contains lines longer than the aforementioned 70 characters per line. While static layouts tend to waste much screen space at larger viewing sizes, more flexible layouts often produce excessively long lines that detract from readability. Particularly text-intensive sites should keep the wide text area at a minimum to allow for efficient reading of longer text in widescreen contexts.
Small text ratio = $\frac{\textit{Small text area}}{\textit{Total text area}}$	Determines how much of the total text area is likely to appear too small on high-resolution displays. There is no such general guideline for what constitutes potentially too small a font on devices with wide-format aspect ratios. Many screen designs have been optimised for a particular resolution, but the small text area should be kept minimal and not substantially increase with higher resolutions.
Visible text ratio = $\frac{\textit{Visible text area}}{\textit{Total text area}}$	Determines how much of the total text area is visible in the viewport. Larger displays should potentially show more text, but too much text at once can overstrain the user. Particularly on the first screen, it is most important to balance out text and media in order to keep the user interested in reading even if continuing requires scrolling.
Visible links ratio = $\frac{\textit{Visible links}}{\textit{Total links}}$	Determines how many of the total navigation options are visible in the viewport. Web sites generally attempt to place as many controls and navigation options as possible in the top navigation area. Larger viewing environments potentially display more links at once and are likely to cause an overload on navigation options.
Media-content ratio = $\frac{\textit{Visible media area}}{\textit{Visible content area}}$	Determines how much of the visible content area is filled with media. Larger viewing sizes should not only fit more text, but scale embedded media accordingly to keep the reading experience interesting. In the case of ads, it may be to the frustration of the user if larger viewports only fill with more ads and to the disadvantage of the advertiser if their ads do not scale with the page's remaining content.

a particular kind occur and the space on the screen that they occupy and use effectively.

The ratios defined in Table 1 directly relate to the problems we have identified in our study. The document-window and content-window ratios assess how efficiently the available screen space is used and how much scrolling is required to view the whole document. The wide text and small text ratios target the overall readability, while the remaining ratios are related to the visibility and spatial distribution of different content elements. The visible text and link ratios hence give further indication of the web site's readability and operability depending on the viewport. Finally, the media-content ratio and possible derivatives track the proportions of text, media and, for example, ads at different viewing sizes.

We note carefully the correlation between the different metrics as optimal results for some of them may have a negative impact on others. For example, higher visible text ratios may be preferable not to impose too much scrolling on users, however the wide and small text ratios should be the constraining factors. The same is true in combination with the media-content ratio as images related to the content will need to be aligned and scaled with the text to keep the reader interested. Finally, an optimal content-window ratio of 1 is a desirable goal in theory, but may come at the expense of aesthetic considerations where the appropriate use of white space can often play an important role to support the overall usability of a web site and improve the user experience.

It follows that optimising web layouts along the defined metrics means to balance out the measured values, as it is impossible to achieve optimal numbers for all considered factors at the same time. Depending on the particular content and purpose of a web site, developers may put more emphasis on some aspects of the layout at the expense of others. The main purpose of online newspapers, for example, is to deliver news content to the screens of their users and support them in accessing this text-heavy information. Special emphasis should hence be on all factors concerning readability on the screen. This may in return mean lower values for the media-content ratio compared to, for example, more artistic web sites where information is mainly conveyed through images and not text. Furthermore, how accurate the obtained values will be depends on how fine-grained the measurements are performed. If we only measure by means of rectangular shapes, it can be rather difficult to approach the effective text area when the actual shape of the text may be significantly different after white space caused by wrapping is removed. Also care has to be taken to avoid content elements being measured twice, for example, due to a media area being embedded in text. We show how we implemented the proposed metrics and discuss the encountered problems in the next section.

IMPLEMENTATION

The proposed metrics were implemented on top of the popular jQuery JavaScript library[10] to ease development and maintain cross-browser compatibility. Using jQuery, the document-window ratio is represented by the product of the computed width and height for *$('body')* divided by the respective product for *$('window')*. We need to mark content elements with an empty CSS class *isContent* so that the content area is the sum of all areas determined for the jQuery selector *$('.isContent')*. Putting the content area into proportion with the window area gives the content-window ratio. Because media can be either embedded using *img*, *object* (e.g. for Flash animations) or the new HTML5 *video* tag, the media-content ratio is computed for *$('img,object,video')* under two conditions. First, one has to make sure that only elements with a non-zero width and height are considered (i.e. they actually fill space somewhere in the page). Second, the element's offset must lie within the current viewport and only the visible part of the bounding box counts towards the area (i.e. only the space they effectively fill in the current viewport is measured). Counting the links in the content area via *$('.isContent a')* in proportion to those that are currently visible yields the visible links ratio.

Considerably more effort has to be expended to compute the various text ratios that we have defined. While there are several difficulties due to the way in which browsers interpret font sizes specified using CSS and although it is not possible to count the characters per line in JavaScript per se, relatively accurate results can still be obtained with the following procedure. Many web sites make use of predefined semantic markup such as *h1* to *h6* for headings and *p* to mark paragraphs of text. As the nesting of these elements is not meaningful and therefore discouraged by the HTML specification, we can measure distinct text areas for the elements returned by *$('h1,h2,h3,h4,h5,h6,p')*. To assess areas of small text, we need to convert the font size of each text element from *px* to *cm* to resolve the physical character height on the screen; however, this requires taking into account the current screen's DPI. It turns out that all major browsers directly work with the system DPI of usually 96, which can be significantly different from the current display's actual DPI. In our tests, a square specified at 10cm by 10cm using CSS actually measures between 7.46cm on a 15.4" notebook with 128 DPI at 1650x1080 and 9.47cm on a 30" screen with 101 DPI at 2560x1600. Our current workaround is to prompt the developer for the screen diagonal (a value which cannot be obtained via the *screen* object in JavaScript). In combination with the current screen's resolution, we can calculate the pixels per inch and extrapolate the effective character height when converting to centimeters. Currently, we consider elements with a font size lower than 0.4cm to be part of the small text area. As there is no general threshold for what constitutes small text, this value has been determined by an informal experiment where three members of our research group were separately asked to select the one paragraph that is still comfortable to read from a number of paragraphs of the same text in various font sizes from 0.1 to 0.6cm. As for wide text areas, we can base our implementation on available research results and currently count in particular paragraph elements whose width is greater than that of a hidden *span* created dynamically at run-time to fill 70 characters using the respective font. In the last step, we remove the summed up area for *$('img,object,video', p)* for each measured paragraph *p* in order to exclude any media that may have been counted towards the text area.

[10] http://jquery.com/

Based on this implementation of the proposed metrics, we have started to develop a visual tool that can be loaded into existing web pages at run-time to perform on-the-fly measurements in the browser as shown in Fig. 7. It then allows the user to select content elements (and their child elements) to be included in the measurements by hovering and double-clicking the respective areas in the web page. That way our implementation does not have to rely on the accuracy of automated methods and designers have full control over fine-grained measurements based on the visual representations of the content elements rather than their hypertext specification. The tool draws an outline around all content elements that were considered in the performed measurements. As illustrated in Fig. 7, we carefully chose visual annotations such as outlines, not borders, and coloured overlays not to break the layout and flow of the original page which is the case with many other evaluation tools. Because the viewing context plays an important role in our evaluations of the web layout, the evaluation method is not only sensitive to changes regarding the visible web area (e.g. resizing of the browser window), but in order to react to potential adaptations of content and presentation, also automatically performs recalculations when content is dynamically added or updated in the web page, e.g. via AJAX menus or tab-based navigation controls that essentially toggle the visibility of page elements.

To compare the results obtained for a number of different viewing sizes, the visual tool also allows individual measurements to be captured and restored by clicking the respective links. For each setting that is captured, a new link will be added, showing the window dimensions, the document-window ratio and the difference to the content-window ratio. For example, the first measurement shows that even a maximised window at 2560x1600 still requires users to scroll down 3.47 pages to view the whole article as large amounts of screen space left and right are filled with background. The +2.21 in this case means that only slightly more than one viewport should be required to fit the actual content.

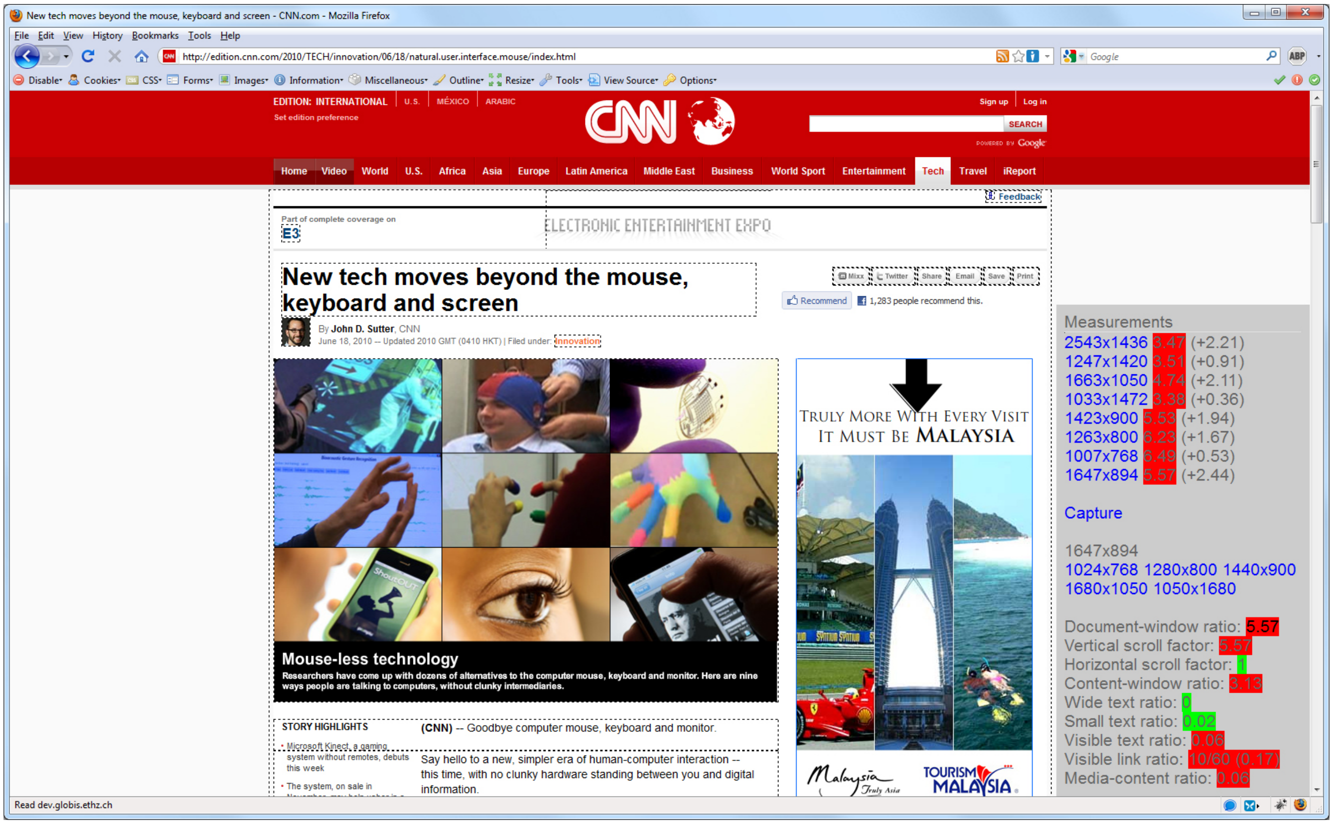

Figure 7. Measuring the quality of a CNN article's layout for a number of different viewing sizes on a 30" screen at a resolution of 2560x1600

The colouring of the obtained values is based on several valuation methods for the different metrics. For example, we say that the document-window ratio is better the closer it is to 1 and expect it to come much closer to this value with increasing viewing sizes. We have defined a number of thresholds for near-optimal, tolerable and poor-quality values based on which the measured values can be assessed and highlighted, for example, in green, yellow and red. However, as the current range of values is largely experimental, we will scrutinise these in another study. The goal at this stage was to show that the metrics we have defined can be implemented in a way that is feasible, with the focus being on the methods required to provide visual feedback for developers so that they can see fairly quickly where adaptations are required. The tool is available for download from our web site[11].

CONCLUSION

In an attempt to set new directions for research in web design where large, high-resolution displays have so far received little attention, we have presented an analysis and comparison of screen usage by text-centric web sites viewed in different widescreen contexts. The results show that the spatial distribution of content elements does not scale well and leads to an increasing amount of unused screen space and no reduction in scrolling. Based on our study, we have proposed several metrics to quantify the critical aspects of screen layouts and have shown how they can be implemented and applied to existing web sites. There are of course many more factors to be considered and directly measuring the quality of web sites in a way that is consistent with how users perceive it is very difficult. For the same reason, it is problematic to define universal guidelines such as ideal values or thresholds and alert levels for the numbers computed by our tool. While we can argue that a small text ratio of 0.02 measured by our tool is worse than one of 0, realistic thresholds and definite alert levels depend on the particular type of web site under development. With the study presented in this paper and the proposed metrics, it will be interesting to correlate any numbers produced by the tool with data collected in user studies in order to scrutinise some of the more experimental and informally determined placeholder thresholds currently used in our implementation of the metrics.

Despite our study being based on a significant number of web sites, the approach could still be criticised for focussing on a specific type of web sites. We intentionally chose online newspapers in order to compare a relatively homogeneous set of web pages and focus on quality aspects that are directly related to the common task that the reviewed web sites were designed for, namely delivering text-centric content to the screens of their readers. Furthermore, we believe that many of our findings can be generalised to news sites and aggregators not considered in the study and potentially to other prominent classes of modern web applications that

[11] http://dev.globis.ethz.ch/jqmetrics

also display large amounts of textual information, such as blogs, wikis and discussion forums.

Our future work will mainly follow two directions. First, we plan to design and implement an adaptive screen layout for news web sites based on the proposed metrics while exploring different forms of context-aware adaptation including the use of multi-column layouts and different scrolling mechanisms. Second, we want to conduct user studies to be able to compare how users actually perceive the quality attributes assessed by our tool and to learn which web site adaptations bring real benefits in widescreen contexts.

REFERENCES

1. A. Borning, R. Lin, and K. Marriott. Constraint-Based Document Layout for the Web. *Multimedia Systems*, 8(3), 2000.
2. C. Braganza, K. Marriott, P. Moulder, M. Wybrow, and T. Dwyer. Scrolling Behaviour with Single- and Multi-column Layout. In *Proc. WWW*, 2009.
3. Y. Chen, W. Ma, and H. Zhang. Detecting Web Page Structure for Adaptive Viewing on Small Form Factor Devices. In *Proc. WWW*, 2003.
4. M. Czerwinski, G. Smith, T. Regan, B. Meyers, G. Robertson, and G. Starkweather. Toward Characterizing the Productivity Benefits of Very Large Displays. In *Proc. INTERACT*, 2003.
5. M. Dyson. How physical text layout affects reading from screen. *Behaviour & Information Technology*, 23(6), 2004.
6. L. Findlater and J. McGrenere. Impact of screen size on performance, awareness, and user satisfaction with adaptive graphical user interfaces. In *Proc. CHI*, 2008.
7. J. H. Goldberg and J. I. Helfman. Evaluating User Expectations for Widescreen Content Layout. In *Proc. UPA*, 2007.
8. J. H. Goldberg, J. I. Helfman, and L. Martin. Information Distance and Orientation in Liquid Layout. In *Proc. CHI*, 2008.
9. J. Hartmann, A. De Angeli, and A. Sutcliffe. Framing the User Experience: Information Biases on Website Quality Judgement. In *Proc. CHI*, 2008.
10. G. Hattori, K. Hoashi, K. Matsumoto, and F. Sugaya. Robust Web Page Segmentation for Mobile Terminal Using Content-Distances and Page Layout Information. In *Proc. WWW*, 2007.
11. E. Insfran and A. Fernandez. A Systematic Review of Usability Evaluation in Web Development. In *Proc. WISE Workshops*, 2008.
12. M. Ivory and R. Megraw. Evolution of Web Site Design Patterns. *ACM Trans. on Information Systems*, 23(4), 2005.
13. M. Ivory, R. Sinha, and M. Hearst. Empirically Validated Web Page Design Metrics. In *Proc. CHI*, 2001.
14. C. Jacobs, W. Li, E. Schrier, D. Bargeron, and D. Salesin. Adaptive Grid-Based Document Layout. *ACM Transactions on Graphics*, 22(3), 2003.
15. T. Lavie and N. Tractinsky. Assessing dimensions of perceived visual aesthetics of web sites. *Int. J. of Human-Computer Studies*, 60(3), 2004.
16. R. Levering and M. Cutler. The Portrait of a Common HTML Web Page. In *Proc. DocEng*, 2006.
17. P. Lynch and S. Horton. *Web Style Guide, 3rd edition: Basic Design Principles for Creating Web Sites*. Yale University Press, 2008.
18. M. Matera, F. Rizzo, and G. Carughi. Web Usability: Principles and Evaluation Methods. *Web Engineering*, 2006.
19. A. Nanavati and R. Bias. Optimal Line Length in Reading–A Literature Review. *Visible Language*, 39(2), 2005.
20. G. Robertson, M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith, and D. Tan. The Large-Display User Experience. *IEEE Computer Graphics and Applications*, 25(4), 2005.
21. E. Schrier, M. Dontcheva, C. Jacobs, G. Wade, and D. Salesin. Adaptive Layout for Dynamically Aggregated Documents. In *Proc. IUI*, 2008.
22. H. Weinreich, H. Obendorf, E. Herder, and M. Mayer. Not Quite the Average: An Empirical Study of Web Use. *ACM Transactions on the Web*, 2(1), 2008.
23. B. Yost, Y. Haciahmetoglu, and C. North. Beyond Visual Acuity: The Perceptual Scalability of Information Visualizations for Large Displays. In *Proc. CHI*, 2007.