Green Patterns in Web Design

Master's Thesis

Dyon van der Ende

Green Patterns in Web Design

THESIS

submitted in partial fulfilment of the
requirements for the degree of

MASTER OF SCIENCE

in

COMPUTER SCIENCE

by

Dyon van der Ende
born in The Hague, the Netherlands

Software Engineering Research Group
Department of Software Technology
Faculty EEMCS, Delft University of Technology
Delft, the Netherlands
www.ewi.tudelft.nl

Green Patterns in Web Design

Author: Dyon van der Ende
Student id: 4712730

Abstract

The internet is a system that was introduced over 30 years ago and has taken over the world since then. While the advantages are unmistakable, we must also acknowledge that the internet is a large contributor to the emission of greenhouse gasses. The footprint of the individual user might be relatively small, but as the number of users is only expected to grow, we must find a way to make the internet more sustainable. In this thesis we explore the state of sustainable web design from three different perspectives: *academic*, *development* and *end user*. Based on academic publications we create a catalogue of nine green web patterns that have empirically been shown to reduce the energy consumption. Secondly, we analyse developer communication on GitHub to see what green patterns are being used in practice. Our results show that there is very limited conversation about the implementation of green patterns, whereas for mobile app development the use of green patterns is more common. Lastly, we take on the end user's perspective and audit the websites of universities to find a relation between a sustainable reputation and a sustainable website. For the audition of websites we create a Google Lighthouse plugin specifically to test sustainability. We find a significant difference between the websites from the group of most sustainable universities and the group of least sustainable universities. The main takeaway of this thesis is that there is a lot of room for improvement to make websites more sustainable. A lack of awareness found from all three perspectives is currently a bottleneck for wider adoption of green patterns. We argue that the availability of development tools with built-in sustainability features could increase awareness and adoption of sustainable web design.

Thesis Committee:

Chair: Prof. Dr. A. van Deursen, Faculty EEMCS, TU Delft
Supervisors: Dr. L. Cruz, Faculty EEMCS, TU Delft
Dr. J. Sallou, Faculty EEMCS, TU Delft
Committee Member: Dr.ir. E. Aivaloglou, Faculty EEMCS, TU Delft

Contents

Chapter 1

Introduction

In March 1989 Tim Berners-Lee wrote a proposal to explain to his colleagues at CERN the benefits of a global hypertext system [7]. His proposal turned out to be the foundation of what we have come to know as the internet, or the web. The hypertext system was introduced as an information management system that could help with structuring the many digital documents that are constantly being created and updated in a large organisation like CERN. Hypertext is an unconstrained way of linking from one document to another, a concept familiar to any user of the internet nowadays as hyperlinks. A year later, in 1990, Berners-Lee started using the name World Wide Web, when he created the first web browser, a hypertext editor that could process HyperText Markup Language (HTML) files (Figure 1.1). Later in 1995 and 1996 JavaScript (JS) and Cascading Style Sheets (CSS) were released, respectively. JavaScript is a programming language designed to run in the browser and can interact with the content of the web page. CSS is a style sheet language used to style the look of a web page, for instance by setting custom colours, fonts and sizes. Together HTML, JavaScript and CSS form the core technologies behind every website on the internet. Nowadays, over 5 billion people worldwide use the internet, of which the majority on a mobile device [23].

1.1 Motivation

The internet as we know it today is a complex series of technologies working together to handle requests and send files between servers and devices over layers of infrastructure. Every action on the internet causes a chain reaction of requests that travels between the communicating devices, using energy at every layer in this network. A user of the internet emits between 28 and 63 grams of CO_2 and uses up to 35 litres of water and 20 cm^2 of land per gigabyte [33]. Andrae expects a total consumption of 3234 TWh in 2030 for the internet as a whole, which in some scenarios could be half of the world's energy capacity [4]. A reckoning from Dornauer and Felderer shows that mobile phone users spent half their time on the internet. A 5% reduction of mobile energy usage is equivalent to closing one Fukushima nuclear power plant [15]. The COVID-19 pandemic has demonstrated how dependent we are on the internet. During lockdowns and stay-at-home orders in Europe online activities increased significantly, with an increment of over 200% for remote working applications [19].

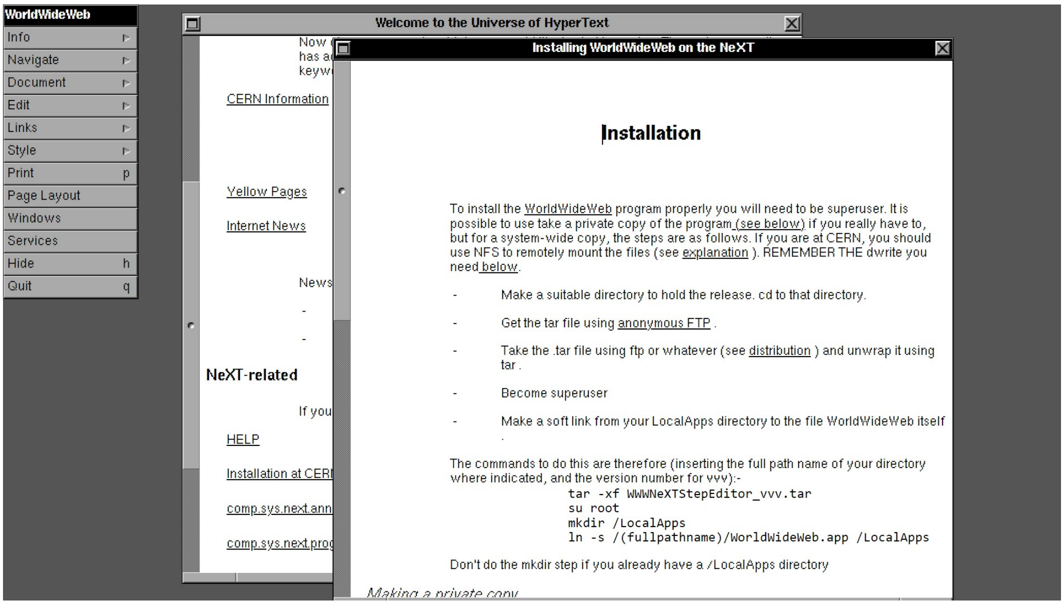

Figure 1.1: CERN hosts an online version of the first WorldWideWeb browser created by Tim Berners-Lee.

According to Itten et al. the internet is an *egalitarian* technology, meaning that the footprint of each individual user is small [26]. Web developers can play a role in reducing the footprint of many individual users at once by making their websites more sustainable. Understanding how to design a sustainable web page is quite challenging due to the complexity of the internet and developers find themselves not well equipped to do this [25]. There are so many variables that need to be considered that even seemingly simple things turn out to be complicated. The rebound effect, or Jevens paradox, is an example of this. This effect occurs when something is made more efficient, which increases the consumption, leading to an overall higher use of resources [40]. The internet has become more efficient over the years as technologies have advanced, but its usage has grown faster than the efficiency improvements could account for. But as global warming is a problem that requires immediate action, the emissions caused by the internet can no longer be ignored.
With the Paris Climate Accords of 2015, 196 countries agreed to achieve a 50% reduction of emissions by 2030. Since then, many companies, institutes and organisations, among which many universities, have committed to this goal or an even higher reduction. The TU Delft, for example, writes the following [41]:

> “TU Delft aims to be carbon neutral [. . .] by 2030. We focus on all aspects that affects our climate and the environment: from buildings and energy systems to waste management and mobility.”

A sustainable website is a necessity to achieve this goal. Nonetheless, according to the Website Carbon Calculator, the TU Delft’s website scores worse than 61% of all webpages globally and emits 0.66 grams of CO_2 per visit.[1] It illustrates how, despite great ambitions, the website has a lot to gain from a more sustainable design.

1.1.1 Sustainable Web Design

Sustainable web design is the discipline that concerns itself with building a more sustainable internet. An important aspect of sustainable web design are design patterns. Design patterns are "general, reusable solutions to recurrent problems in software design [12]." In the case of sustainable web design, the problem is the inefficient consumption of energy of a website. Accordingly, the solutions are targeted at reducing energy consumption. Green patterns are design patterns that are focused on energy consumption and can be considered as guidelines on how to make designs more sustainable.
Researchers have a significant role in exploring and verifying green patterns. They can find more efficient ways to make use of existing technology, or create new technologies that are an improvement over the old technologies.
Developers can then use the green patterns in their design to make websites more sustainable. Design patterns are valuable to developers as they make it easier to design better websites, without having to reinvent the wheel. They can trust that the patterns have been tested and have the desired effect.
Finally, the websites are used by the end user, sometimes referred to as the *client*. In general, the end user has no technical background and just expects a website to work. For someone with no technical background it can be hard to verify whether a website is sustainable or not. It is likely that the user will associate the sustainability reputation of a brand with the sustainability of its website. Whether such a relation between reputation and sustainability exists is unknown as this has not been investigated before.

1.2 Research Questions

In this thesis we set out to explore the state of sustainable web design. We define three perspectives to explore the subject from as they provide a good overview. The perspectives are: *academic*, *development*, and *end user*. For each perspective a research question is formulated:

RQ1 **What are green web design patterns supported by literature?**

RQ2 **Which design patterns do web developers adopt to improve energy efficiency?**

RQ3 **Is there a relation between the sustainability of a university and the sustainability of their website?**

For the first research question we study articles that investigate the relation between a design pattern and energy consumption. Patterns that have been empirically proven to reduce energy consumption are included in a catalogue of green web patterns. In total we find nine patterns that are added to the catalogue.

[1] `https://www.websitecarbon.com/website/tudelft-nl/`

For the second research question we mine GitHub repositories related to web development. By analysing the conversations in commit messages, issues and pull requests we determine which green patterns are used in practice by developers. Our data suggests that the implementation of green patterns currently is very limited, as less than one percent of the analysed discourse mentions any form of energy optimisations.
For the third research question we compare websites of the highest and lowest ranked universities on the QS World University Rankings to each other [34]. With a sustainability plugin for Google Lighthouse we audit to what extent the universities have a sustainable website. Statistical analysis shows that there is a significant difference between the two groups to the advantage of the higher ranked universities.

1.3 Contributions

The following contributions are made:

- a catalogue of green web patterns found in the literature;
- an analysis of green patterns used by web developers on GitHub;
- an initial implementation of the Web Sustainability Guidelines in the form of a Google Lighthouse sustainability plugin;
- an analysis of the relation between reputation and sustainability of 200 university websites;
- all code and data are publicly available for further research in the form of a replication package.

1.4 Thesis Outline

The remainder of this thesis is structured as follows:

- chapter 2 discusses related works;
- chapter 3 answers **RQ1** by creating a catalogue of green web patterns;
- chapter 4 answers **RQ2** by mining GitHub repositories;
- chapter 5 answers **RQ3** by comparing the websites of universities with a Google Lighthouse plugin;
- chapter 6 discusses how the internet can be made more sustainable based on the findings of the previous chapters;
- the thesis ends with a conclusion in chapter 7.

Chapter 2

Related Work

Web scientist can contribute to a more sustainable future by designing their products in a way that reduces the power consumption and social and economic costs during the life cycle of their products, i.e. mitigating the negative impact [31]. Since sustainability can be hard to quantify, we will use this chapter to demonstrate how it is studied in the literature. First, we give an overview of methods that are used to measure sustainability in the field of sustainable web design. Then we list approaches to improve web sustainability, known as green patterns. We end this chapter with a section on the awareness of developers regarding green web patterns and a section on non-academic work. For a description of how relevant works were collected we refer to section 3.1.

2.1 Measuring Sustainability

Using the definition from Greenwood, sustainable web design is “an approach to designing web services that prioritizes the health of our home planet [21].” Since this is hard to quantify, generally energy consumption and CO2 emissions or CO2-equivalent are used as measurements. Sometimes the land and water usage of data centres is included as well.

In the cases where we want to measure the energy consumption, an energy profiler (software based) or a power monitor (hardware based) can give a good idea of the power consumption of the software on a single device [11, Chapter 2]. More specifically for websites, RECON was created [8]. It is a machine learning model that can estimate the energy consumption of a mobile web page within an error of 7%. For more detailed insights it also provides information on specific components like HTML and JavaScript. The advantage of RECON is that it takes the website as input and gives an estimation, eliminating the need to monitor the entire host system. Small updates to a design can easily be tested on their efficiency with this method.

Since 2022 the Firefox web browser ships with an energy profiler included in the developer tools. It can display estimated power draw of system components like the CPU and the GPU and shows the CO_2-equivalent, see Figure 2.1. With support for Linux, Mac and Windows, this tool is one of the most versatile solutions for developers currently available.

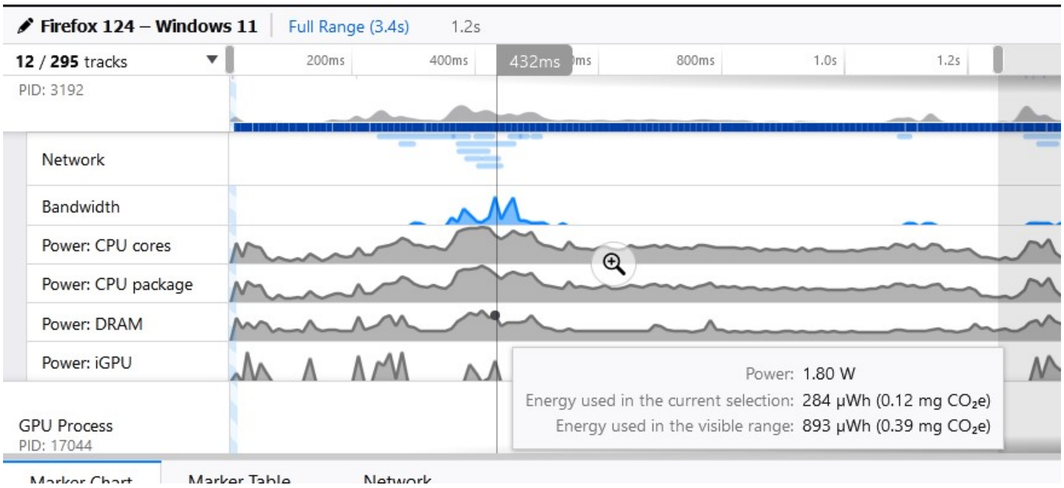

Figure 2.1: Firefox comes with an energy profiler in the developer tools. This screenshot shows the energy usage of loading a website.

2.2 Developer Awareness

Since developers play a key role in making the web more sustainable, there have been several surveys that investigate awareness and struggles that developers face in making this a reality. In-depth interviews with 6 developers working in the industry revealed unawareness of energy patterns in a study from Rani et al. [35]. Despite this unawareness, the developers showed concern for anti-patterns that, instead of reducing energy consumption, lead to higher power consumption.

Hulleberg et al. investigated the awareness of developers on the impact of the web on sustainability [25]. A questionnaire consisting of 33 questions explored several aspects of sustainable web design, related to images, videos, fonts and compression. On average the 77 participants of the online survey answered 69% of the questions correctly. The authors conclude that there is a discrepancy between awareness and implementation of green practices regarding sustainable web design.

Hansen et al. interviewed 27 Norwegian web developers who were either practitioners, prospective practitioners or worked in academia [22]. Through a thematic analysis they identified several barriers. The first and second barrier are caused by a knowledge gap, where developers expressed that they did not feel well equipped to implement sustainable solutions and do not comprehend how the internet affects sustainability. There is a lack of good resources and the ones that are currently available are not easily understandable. If good guidelines would be available, they could help remove the barriers.

A third barrier that was identified was the customer-oriented prioritisation. The customers for whom the developers make the products often prioritise accessibility and code maintainability over sustainability. This is partly due to the fact that there are no regulations that require sustainable web design, whereas this is the case for accessibility and maintainability. The Web Content Accessibility Guidelines are often cited in the interviews, as Norwegian law enforces accessibility in web design for public and private organisations.

The fourth barrier is the lack of motivation within the industry. Some interviewees said that their colleagues are not easily convinced of the importance of sustainable web design. Even though most of the participants showed an interest in it to some extent, 63% of them did not give it priority in their decision making process.
Lastly, in education sustainable web design is not part of the curriculum. The lack of sustainability in academia is something that universities and other educational institutes could resolve by offering courses around sustainability in their programmes.

2.3 Green Web Patterns

Battery life is one of the main reasons for research interest in energy optimisations for web applications. The focus in these works is therefore on mobile web and not so much on desktop. A survey about energy saving strategies on mobile web categorised the different approached in three groups: adoption of web content, saving processor utilisation and reduction of web traffic [15].

Adoption of web content The paper “Who killed my battery? Analyzing mobile browser energy consumption” illustrates the importance of battery life in research [39]. In this work mobile web pages are broken down into the components *images, CSS, JavaScript* and *others* to see how these affect the energy usage. A 30% reduction of the energy consumption was shown for the mobile Wikipedia site by removing unused JavaScript code. A similar reduction was found when shrinking the CSS files on Apple’s website, at about 40%. The last optimisation is to use the JPEG format for images instead of PNG, BMP and GIF. As part of this study the authors also investigated off-loading compute intensive tasks like image rendering via a proxy server and via a back-end server. The idea behind this makes sense from the perspective of saving battery life on the client device. Google Flywheel [2] and WEEC [44] are other examples of intermediate servers that optimise websites before sending it to the user. Although these solutions also had a positive impact on the energy consumption of the user’s device, the overall benefit of this approach on sustainability is questionable. By off-loading tasks the computations are not reduced but simply relocated to a different device. Theoretically, that device could still be more efficient, but then the overhead of the extra communication has to be taken into account as well.
Displays are known to be a large consumer of energy on smartphones. Using the characteristics of OLED displays Li et al. found that using dark colours in the design of a website could lead to a 40% reduction in the power consumption [29]. Nyx is a tool developed to exploit this and redesigns a web page from the original colour scheme to one that uses dark colours, while preserving the relation between the original colours. From the evaluation it appears that 97% of users are open to using Nyx to preserve power when their battery is running empty [30].

Saving processor utilisation GreenWeb takes a unique approach to developing energy aware web applications [48]. It is a language extension for CSS that allows developers to target a quality-of-service experience. Interactions on a web page come with certain expectations, for example touching an object should be processed fast, but a loading screen can take longer. By annotating this, the runtime can decide which parts of the site to optimise, while delivering an acceptable amount of performance. The benefits of the extension are significant with energy savings up to 66%. To assist developers with annotating the tool AutoGreen can automate the process. Despite this, GreenWeb has not seen any updates since its introduction in 2016 and there are no indications that it is used in practice.

Reduction of web traffic Although the more efficient file formats and minification mentioned earlier technically lead to a reduction of web traffic as well, Dornauer and Felderer use this category for methods that use caching and (pre-)filtering of content [15]. Caching can save data communication by storing static elements on a device so that they don't have the be downloaded a second time, which can save a significant amount of energy [18]. How large these savings can be was illustrated by Albasir and Naik, who developed a method to reduce the number of ads shown on a mobile page based on the battery level [3]. By decreasing the number of ads 30% less battery was used and 40% less bandwidth.

2.4 Repository Mining

A common approach for finding green patterns is through repository mining, a method that we use in chapter 4 as well. Data is extracted and filtered from software repositories based on programming patterns and selected key words. A thematic analysis on the filtered data can reveal strategies to achieve energy optimisations. Repository mining is a useful method to gain insight into what green patterns are applied in practice by developers and it can be applied to platforms other than the web too.

Moura et al. used GitHubArchieve to analyse roughly 96 million commit messages looking for energy aware commits [32]. Commits were selected when matching the terms **energy consum*, *energy efficien*, *energy sav*, *save energy*, *power consum*, *power efficien*, *power sav** and **save power**. These were then categorised based on intention and code implementation. 371 commits were labelled as energy-aware and from a thematic analysis the authors concluded that authors mostly use low-level optimisations, such as frequency scaling.

Bao et al. extended this study, but with a focus on open-source Android applications [5]. They found 468 power management commits from 154 apps, using a less restrictive set of keywords: *power, energy, battery* and *wakelock*. About 12% of all included apps used power management solutions.

Cruz and Abreu analysed the energy related issues, pull requests and commit messages of GitHub repositories of open-source Android and iOS apps [12]. The keywords used in this study are *power, energy* and *battery*, dropping the *wakelock* keyword from the previous study. Through a thematic analysis of the data a catalogue of 22 green patterns for mobile app development is formulated.

2.5 Non-academic initiatives

From outside of the academic world there are many initiatives for a greener web that are worth mentioning. SustainableWWW[1], Sustainable Web Manifesto[2] and The Green Web Foundation[3] are all organisations with the goal of making the internet more sustainable. *Branch* is an online magazine by ClimateAction.tech about internet sustainability and is hosted on a website that changes its appearance based on the availability of clean energy on the grid.[4] Carbon calculators are available online that estimate a website's emissions based on CO2.JS[5], a model developed by The Green Web Foundation that is also used in Firefox's energy profiler. And besides the W3C's Web Sustainability Guidelines [13] that will be discussed in chapter 5, several guidelines have been developed. An open-source project on GitHub collected 115 web ecodesign best practices[6] and The European Institute for Sustainable IT has the Handbook of Sustainable Design of Digital Services with 61 recommendations[7].

[1] http://sustainablewww.org/
[2] https://www.sustainablewebmanifesto.com/
[3] https://www.thegreenwebfoundation.org/
[4] https://branch.climateaction.tech/
[5] https://www.thegreenwebfoundation.org/co2-js/
[6] https://github.com/cnumr/best-practices/
[7] https://gr491.isit-europe.org/en/

Chapter 3

A Catalogue of Green Web Patterns

During the literature study phase of this research, it was noticed that several papers mentioned the absence of guidelines for sustainable web design [43, 6, 22, 25, 35]. At the same time, we also noticed that often green patterns are introduced in articles without a source. We have therefore decided to list all of these techniques found during our research and present them as a catalogue of green web patterns. The catalogue can be used by researchers as a resource or used by developers as a guideline. With this catalogue we answer research question **RQ1: what are green web design patterns supported by literature?**

This chapter is structured as follows: in section 3.1 we describe the methodology for the creation of the catalogue, followed by the catalogue itself in section 3.2. The chapter ends with a discussion in section 3.3 and a conclusion in section 3.4.

3.1 Methodology

For this thesis the collection of related works was performed through Google Scholar. We did not follow a systematic procedure, but worked with an initial collection of articles that kept on being expanded as the research progressed and as new papers were published. By using combinations of the following key terms as search queries the initial set of articles was constructed: *(sustainable / green / energy / efficiency / battery)* + *(web / internet / progressive web apps)*. Based on title and abstract it was decided if the articles were potentially relevant. This collection was expanded by adding cited works as well. In the research phase for chapter 5 more targeted search queries were added like *code minification*, *web fonts*, *page weight*, *image format* and *video codec efficiency*.

For the construction of the catalogue of green web patterns we followed a strict inclusion requirement. Patterns are included in the catalogue only if the corresponding paper explicitly mentions an effect of the pattern and energy consumption through empirical testing and can quantify it. Some articles tested several optimisations at the same time which made it difficult to tell the effect of a single optimisation and these are thus not included. In total 12 articles contained patterns that fulfilled the requirement.

Table 3.1: Green web patterns found in the literature

	Pattern	Source
1	Use modern font formats	[17]
2	Optimise Lighthouse performance score	[9]
3	Consider the use of WebAssembly	[14, 45]
4	Combine external files	[43]
5	Minify code	[43]
6	Offer a dark theme	[29, 46]
7	Scale images	[44]
8	Use modern image formats	[16, 24, 39]
9	Use modern video encodings	[42]

3.2 The Catalogue

We have collected nine green web patterns from the literature. All of these patterns have empirically been shown to have led to a reduction in energy consumption when implemented correctly. In most cases the patterns also have a positive impact on performance and do not have negative impact on the user experience. The patterns are found in Table 3.1 and are described in this section.

3.2.1 Use Modern Font Formats

Custom fonts to style a page are common on the web as 84% of all sites do this [23]. Most fonts need to be downloaded as they are not installed on the users device. The downloaded font file can have several formats: OpenType, TrueType, WOFF and WOFF2. The chosen format impacts the resource usage on the device that it is loaded on. Dornauer et al. evaluates the impact of font formats on CPU, memory, load time and energy and found that WOFF2 is the best format overall [17]. WOFF2 is the most modern file format of the four and the successor of WOFF. It is optimised for web and compression is automatically applied, which gives it an advantage to the others. The other file formats should only be used for fallback, in case the WOFF2-file fails to load. In general, the following order for fallback based on energy consumption is suggested: WOFF2, WOFF, OpenType, and TrueType.

3.2.2 Optimise Lighthouse Performance Score

Intuitively, it makes sense that optimising the performance of a website can lead to better energy efficiency. To test if this intuition holds, Chan-Jong-Chu et al. ran an experiment that tested the correlation between the performance score in Google Lighthouse and energy consumption on mobile Android devices [9]. The results showed a positive correlation, meaning that higher performance scores generally had a lower energy consumption. Especially websites with a performance score of 44 and lower performed significantly worse in energy efficiency.

3.2.3 Consider the Use of WebAssembly

JavaScript is one of the core technologies of the internet, but it is not ideal to handle very compute intensive task. WebAssembly is designed to run these kinds of tasks and is one of the newer additions to the W3C standard being released in 2017. Programs written in C or other language can be compiled into binary WebAssembly code that runs in the browser. Compiled binary code offers many performance benefits compared to running code through an interpreter as is the case with JavaScript. Besides the performance benefit, WebAssembly also offers significant improvements in energy efficiency of at least 30% [45, 14]. Although for many simpler tasks JavaScript is still a good option, WebAssembly is the better solution for more advanced project.

3.2.4 Combine External Files

For most websites a combination of HTML, CSS and JavaScript is used. The HTML file is the main file that is loaded and all other files have to be loaded with requests. Every request takes time and energy to process and reducing the number of requests is therefore commonly considered a good practice. But as Ünlü and Yesilada point out, empirical evidence that this would also save energy was lacking [43]. In their study they put this practice to the test by combining several CSS and JavaScript files into one file for all CSS code and one file for all JavaScript code. Via a proxy server that concatenated all files they found that this technique could save around 4% of power on mobile devices and even more on laptops.

3.2.5 Minify Code

Souders describes minification as "the process of stripping unneeded characters (comments, tabs, new lines, extra white space, and so on) from the code [37]." Some tools even replace variable names with shorter ones to save more data. File size is typically reduced by minification, so it uses less bandwidth to load and could therefore improve energy consumption. Testing this theory showed that applying minification saves 2.5% of power on mobile, translating to an improvement in battery life of 7% [43].

3.2.6 Offer a Dark Mode

Dark themes have the advantage of a lower power usage on screens with OLED technology. These types of screen can turn individual pixels on or off and do not need a backlight for the whole screen, as is the case with traditional LCD. This means that for displaying black or dark colours less energy is needed. Offering a dark design uses this to save battery on mobile devices, where this screen technology is quite common. A study from 2014 claimed a reduction of 40% in the power consumption of the display [29]. A newer study from 2021 found 12% battery savings [46].

3.2.7 Scale Images

Images make up a large portion of the internet with 99.9% of the pages containing at least one image [23]. Images generally take up more space than HTML, CSS, JavaScript and fonts combined [23]. A simple solution to reduce energy consumption is to serve images that are appropriately scaled for the client. There is no need to serve a high resolution picture if it will be displayed on a mobile device with a small display. If scaling is done before transferring the file – server side instead of client side – it could save not only bandwidth, but also processing power. Substantial improvements can be achieved with this technique, averaging at 24.6% reduction in energy consumption [44].

3.2.8 Use Modern Image Formats

Another good way to save energy is by using modern image formats. A 2012 study showed that JPEG was the best image format compared to PNG and GIF, which were the most common at the time. Since then, newer formats have been released. WebP was released in 2018 and AVIF in 2019, both promising better efficiency than JPEG. Dornauer and Felderer and Hu et al. have shown that it is indeed true that the newer file formats perform better and save more energy [16, 24].

3.2.9 Use Modern Video Encodings

Use of video on the web has seen a modest growth over the last few years, to the point that currently one in twenty pages embeds video [23]. Whereas with images the file format plays an important role in the file size and picture quality, this works different for video. For videos both a video and an audio stream are stored and the video format (e.g. MP4) is the container for these two media. Rather, the codec (e.g. H.264) used to encode the video stream is what matters [25]. The best encodings with regard to energy efficiency are VP9 and H.265/HEVC [42].

3.3 Discussion

There are many blogs available online with patterns on how to make websites more sustainable. Although the patterns are often similar to the green patterns from this chapter, references are commonly missing. It is therefore hard to check whether the effect of the patterns is as promised. A methodology is most of the time absent as well, making it unclear why certain patterns are included and others are not. If the author did not include a certain pattern, is it because it is not actually sustainable, or because of another reason? By describing the methodology with inclusion criteria and citing sources, the catalogue presented in this chapter takes away those uncertainties.

The patterns listed in the catalogue only reflect a small portion of the many possibilities in which a website can be made more energy efficient. By using a strict inclusion requirement we had to exclude many potentially good patterns, as no relation with energy efficiency was proven. There is a general idea that best practices automatically lead to better performance and thus better energy efficiency. That this is not always true is shown by Rani et al. who found that a common pattern, Dynamic Retry Delay, did not make a meaningful difference in energy consumption [35]. This shows the importance of having to actually verify patterns before instead of relying on assumptions.

3.4 Conclusion

By analysing relevant literature we created a catalogue of nine green patterns to answer **RQ1: what are green web design patterns supported by literature?** Web developers can use the catalogue to make more sustainable websites without sacrifices in performance or user experience. For research, the catalogue can function as a foundation for future research and we invite researchers to expand the catalogue with new green patterns.

Chapter 4

Repository Mining for Green Web Patterns

Web developers have a instrumental role in implementing sustainable solutions on websites. Repository mining is a well suited method to gain insight in how developers do this. On GitHub, developers can discuss code changes in issues, pull requests and commits. By analysing the discourse we can find common green patterns that are being used in practice. Previous work on GitHub [32], Android [5] and Android and iOS [12] have used repository mining to investigate the implementation of green patterns, as discussed in section 2.4. In this chapter, we adapt Cruz and Abreu's methodology and answer research question **RQ2: Which design patterns do web developers adopt to improve energy efficiency?** The code and data from this Chapter is made publicly available.[1]

This chapter is structured as follows: section 4.1 describes the methodology of the experiment; results are presented in section 4.2; the discussion is found in section 4.3; and the chapter ends with the conclusion in section 4.4.

4.1 Methodology

For this experiment we analysed the issues, pull requests and commits of 2000 code repositories on GitHub, a sample size comparable to Cruz and Abreu [12]. To find relevant repositories we used the search functionality available through the GitHub API with the search query *website*. We used the default parameters for the search function, which ranks results based on relevance, but how exactly this ranking algorithm on GitHub works is undisclosed.

The API has a strict rate limit on the number of request that would have taken an unrealistic amount of time for this experiment to be completed. In order to deal with these restrictions we had to rely on third party solutions PyDriller and The Stack. PyDriller is a Python framework developed to support research in Mining Software Repositories [38]. The tool allows to extract information about commits, developers, modified files, diffs, and source code from repositories.

[1] `https://github.com/dyonende/GreenWebPatterns`

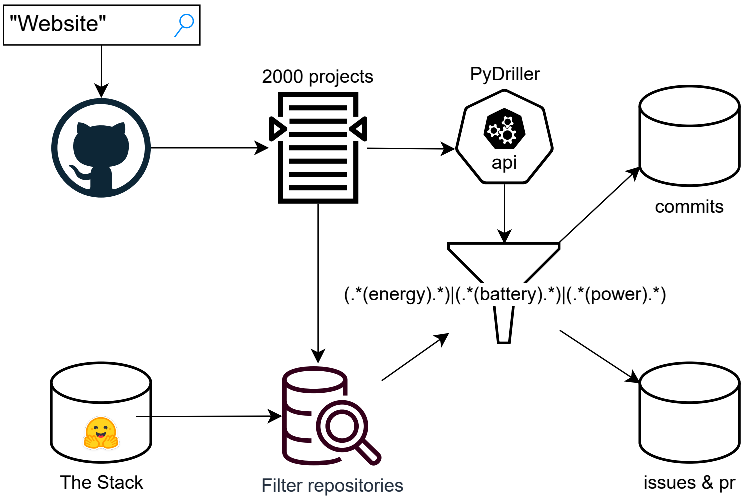

Figure 4.1: Schematic overview of the data collection process.

Since PyDriller works for commits, but not for issues and pull requests, we had to resort to a different method to collect this information. As alternative we used The Stack, a dataset of over 6TB of permissively-licensed source code files, available on Hugging Face [28]. This project collected a large amount of GitHub data for machine learning purposes and as part of it they have also stored issues and pull requests. In total this collection contains 30.9 million rows of data from repositories created between January 1st, 2015 and March 31st, 2022. From The Stack the same set of repositories was selected and the issues and pull requests are filtered.

After these steps the data was collected and filtered with a regular expression to only keep the texts relevant to energy efficiency. The keywords used for this are the same as from Cruz and Abreu and designed to keep messages that contain the words *energy*, *battery* or *power*. In Figure 4.1 the setup of the data collection is shown.

Finally, after the data collection, every issue, commit and pull request was reviewed manually to determine actual relevance. It is possible that messages containing the keywords do not discus energy efficiency and are thus not relevant. In the cases where a text is relevant, it needs to be determined what energy pattern is discussed, so that these can be clustered. Relevance is judged based on the whole text, so it is possible that the part of the text that matches with the keywords is not the part that is relevant.

4.2 Results

Through PyDriller 1119 commit messages from the 2000 repositories were collected after filtration of keywords; from The Stack 1097 were retrieved. Commits generally are shorter messages that describe the change made to code, whereas issues and pull requests are longer conversations between different people that discus ideas or new features. In total only 6 relevant text items were found in the commits, issues and pull requests.

Many of the commit messages that matched with the key words where false positives, for example the phrase *powered by* is used very often. There were also many instances where energy consumption optimisations are mentioned, but this was almost always in the content of a web page. For example, the Ethereum website's repository has a page explaining the energy efficiency of its blockchain system. Something similar applies to documentation pages for desktop applications, where energy consumption is mentioned often.

Comparing the results to the earlier mentioned studies, our findings are dissimilar, see Table 4.1. Their results have at least 17% of the data containing green patterns, whereas in our data less than one percent of the texts does.

Table 4.1: A comparison of the results to the outcomes of similar studies.

	GH [32]	**Android** [5]	**Android + iOS** [12]	**Web**
# keyword match	2189	1457	6028	2216
# relevant text	371 (17%)	468 (32%)	1563 (26%)	6 ($<$1%)

4.3 Discussion

There are several explanations as to why we see such different results compared to the studies with a similar methodology.

One reason could be that the analysed repositories are not a good representation of web development. Some of the repositories included in the data were frameworks for building websites, such as Hugo. Another large group was the documentation pages for software projects like Kubernetes. These documentation pages gave many false positive hits because of changes to the content of the documentation and not the energy consumption of the website. We have tried different search queries, such as *Progressive Web Apps*, but found that these did not give the desired results.

A second reason could be that our keywords are suitable for native application development, but not for web development. Web development is generally higher level than native apps and it is possible that a different set of keywords would have given better results.

A third possibility is that during the manual labelling of the data relevant texts have been overlooked. Labelling is a difficult task and it required roughly 1 hour per 100 texts. It requires expert knowledge on the programming languages and frameworks discussed in the data. Context is often lacking because texts were analysed without any other information like existing features or previous code changes.

However, even if some relevant texts were missed, the number of relevant texts is still low. Based on surveys among web developers, we know that priority and awareness on sustainable web design and energy patterns is low [35, 25, 22]. The results of this experiment are in line with those surveys.

4.4 Conclusion

In this chapter we have analysed 2000 GitHub repositories related to web development to answer **RQ2: Which design patterns do web developers adopt to improve energy efficiency?** Of all the issues, pull requests and commits from the repositories less than 1 percent discusses energy related optimisations. Given that for other platforms the percentages are substantially higher, the discrepancy could indicate a lack of awareness or priority among web developers. We conclude that currently the use of green patterns by web developers is very limited.

Chapter 5

Sustainability of University Websites

It can be difficult to know how sustainable a website is. Many organisations want to give the impression of being very sustainable, while this is not always the case. The sustainability of a website can be determined through tools like a carbon calculator, a green hosting checker or with an energy profiler. Carbon calculators estimate emissions and return a CO_2-equivalent per page view. Green hosting checkers look if the host provider runs on renewable energy. Energy profilers measure the energy consumed on the device while loading and using a web site. While all three methods are useful, they provide limited insight into what makes a website actually sustainable and what could be improved. To get this kind of feedback you would need to inspect the website under the hood with tools like Google Lighthouse. Lighthouse inspects a website with audits in multiple categories: *performance*, *accessibility*, *progressive web app* and *search engine optimisation*. It provides detailed feedback of what can be improved and the potential benefits. For each category it combines the outcome of the individual audits into a single score. A correlation between the performance score and energy consumption was found by Chan-Jong-Chu et al. [9].
Universities have been subject to research in web sustainability in the past. A 2023 study analysed universities in France and Italy with a carbon calculator and found that there are major differences in energy consumption, with some of the worst websites consuming up to 70 times more energy [6]. An article from 2017 showed that the Universiti Malaysia Sabah's website had a lot to gain from optimisations such as compression and minification [36]. In chapter 1 we showed that the TU Delft website is relatively unsustainable with 61% of all webpages having a lower footprint.
In this chapter we will also take university websites as subject for our study and answer research question **RQ3: Is there a relation between the sustainability of a university and the sustainability of their website?** For this purpose we develop a plugin for Google Lighthouse that adds a new category for sustainability. We then use this plugin to analyse the websites of universities. Based on their position on the QS sustainability university rankings 2024 we compare the 100 highest ranked and the 100 lowest ranked universities [34].
This chapter is structured as follows: in section 5.1 the development of the Lighthouse plugin is explained; section 5.2 describes the methodology of the experiment; results are presented in section 5.3; the discussion is found in section 5.4; and the chapter ends with the conclusion in section 5.5.

5.1 Lighthouse Plugin

Google Lighthouse makes it easy to develop plugins by providing documentation and example code. The Green Web Foundation has already created a plugin that audits if a page is built using resources from servers running on green energy. For building a sustainability plugin, a set of audits needs to be selected and implemented that can provide useful insights of the sustainability of a website. As a starting point the Web Sustainability Guidelines from the W3C Sustainable Web Design Community Group are used [13]. The Web Sustainability Guidelines provide a set of guidelines that provide a short description, formulates success criteria and gives examples and resources. This makes it one of the most thorough and complete set of guidelines currently available and especially the success criteria are useful for implementation. With 93 guidelines it is not feasible to include all of them in this research. Most of the guidelines also require some sort of internal access to the back-end or the design process that is not possible with Lighthouse audits. A selection of the guidelines to be implemented has been made based on two criteria: 1. The guideline can be checked through static analysis or can be automated; 2. The guideline is supported by academic literature.

Based on the first criterion 5 guidelines were selected [10] with 7 success criteria, of which 2 could not be backed up by the literature. These guidelines with their success criteria and supporting literature are found in Table 5.1 and cover typefaces, minification, expected files, plain text files and compression. In Table 5.2 the implementation of the audits based on the success criteria are listed. For some of the guidelines there already exists an audits in Lighthouse that can be imported in the plugin, which is indicated in the table. Each audit returns a score between 0 (bad) and 1 (good) based on the results of the audit. The individual scores are weighted and make up the final score. Weights are derived from data from the Web Almanac [23, Chapter 21].

The plugin is developed open-source and additional audits can be added in the future.[1]

5.2 Methodology

For this experiment the QS World University Rankings 2024 were taken as a starting point. This ranking contains 1397 universities from 95 countries and ranks them based on social impact, environmental impact and governance to form a sustainability score. Place 1 to 700 are ranked individually, whereas lower placed universities are put in brackets. Since we are interested in comparing the sustainability of a university to the sustainability of their website we selected two groups: the Top 100 (ranked) and the Bottom 100 (unranked, randomly selected from the 199 universities in the bracket 1201+).

For each university we used Lighthouse with the sustainability plugin to audit the home page. We did this with the default settings of Lighthouse, which tests only the mobile page and not the desktop variant. The code to run this experiment is made public available.[2]

[1] `https://github.com/dyonende/lighthouse-plugin-sustainability`

[2] `https://github.com/dyonende/GreenWebPatterns`

Table 5.1: Guidelines from the Web Sustainability Guidelines that can be automatically checked [13]. Guidelines are defined by success criteria, but not all are supported by scientific literature.

Guideline		Success Criterion	Source
2.18	Take a More Sustainable Approach To Typefaces	Use standard system-level (web-safe / pre-installed) fonts as much as possible.	[17]
		Ensure the number of fonts, and the variants within typefaces (such as weight and characters) are limited within a project, using the most performant file format available.	[17]
3.2	Minify Your HTML, CSS, And JavaScript	All source code is minified upon compilation (including inline code).	[43, 2]
3.18	Include Files That Are Automatically Expected	Take advantage of the favicon.ico, robots.txt, opensearch.xml, site.webmanifest, and sitemap.xml documents.	
3.19	Use Plaintext Formats When Appropriate	Utilize standards such as ads.txt, carbon.txt, humans.txt, security.txt and robots.txt.	
4.3	Compress Your Files	If using a CMS, install an applicable plugin to enable on-the-fly server-side compression, such as Brotli or GZIP. Otherwise, use the provided server configuration files to include and tweak the performance-related features to the requirements.	[2]
		Compress your various images, fonts, audio, and video; by reducing the quality and offering different resolutions / dimensions (sizes) before uploading to a server or content management system.	[44, 24] [25, 42] [39, 2, 27]

Table 5.2: The audits in the newly developed Sustainability category for Lighthouse are based on the Web Sustainability Guidelines success criteria. Some audits could directly be imported from other categories in Lighthouse and are indicated with a ✓.

Guideline	Audit	Description	Imported	Weight
2.18	font-family	Checks if web safe fonts are used		10
	font-format	Checks if WOFF2 is used		10
3.2	unminified-html	Minifying HTML files can reduce network payload sizes		1
	unminified-css	Minifying CSS files can reduce network payload sizes	✓	2
	unminified-js	Minifying JS files can reduce network payload sizes	✓	4
4.3	video-codec	Checks if HEVC, H.264, VP9, AV1 are used as codecs		20
	responsive-images	Checks to see if the images used on the page are larger than their display sizes	✓	8
	optimised-images	Determines if the images used are sufficiently larger than JPEG compressed images without metadata at quality 85	✓	8
	text-compression	Ensure that resources loaded with gzip/br/deflate compression	✓	1

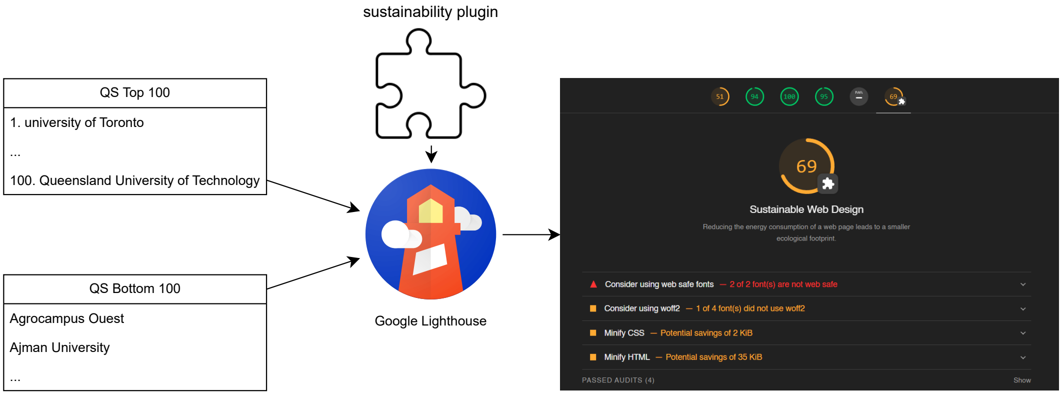

Figure 5.1: A schematic overview of the methodology of the experiment. Two groups of university websites are analysed through Google Lighthouse with the Sustainability plugin. For each university's website a detailed report is produced.

5.3 Results

Of the in total 200 inspected websites there were 8 and 19 failures for the Top 100 and Bottom 100, respectively. The 173 successful analyses each produced a report with results per audit and per category.

5.3.1 Sustainability Score

For the sustainability score the results are plotted in Figure 5.2a. The highest, lowest and average value are all higher for the Top 100. Its maximum value tops out at 100, while for the Bottom 100 the maximum is 79. The difference in mean is 6.8 with 53.1 for the Top 100 compared to 46.3 for the Bottom 100.
To test if there is a significant difference between the two groups for the sustainability score we use the Student's t-test with $\alpha = 0.05$ and the hypotheses

$$H_0 : \mu_{Top100} = \mu_{Bottom100}$$
$$H_A : \mu_{Top100} \neq \mu_{Bottom100}$$

The outcome of the test gives a p-value of $p = 0.0066$ and thus we can reject the null-hypothesis, meaning that there is a significant difference between the two groups.
If we zoom in at the Top 100 and plot the ranking of the university to the results from Lighthouse, the trend line shows only a very slight incline (Figure 5.3a). The Pearson Correlation Coefficient is $r = -0.0774$. The negative correlation is explainable as the most sustainable university corresponds to ranking 1 with the number going up as the sustainability decreases. However, the coefficients are too low to speak of a strong correlation.

5.3.2 Performance Score

The performance score is another interesting metric to look at, because of the earlier mentioned correlation with energy consumption. The distribution of the scores for both the Top

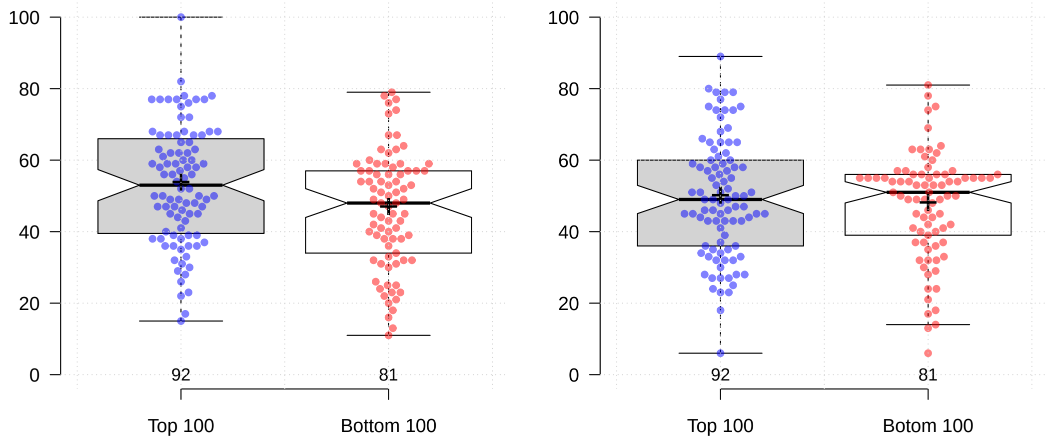

(a) Sustainability score (b) Performance score

Figure 5.2: Lighthouse results for the sustainability score and the performance score.

100 and Bottom 100 in Figure 5.2b shows that the differences are smaller compared to the sustainability score. The Top 100 has a higher mean at 49.4, which is 2 points higher than the 51.4 of Bottom 100, but a lower median (51 to 49). However, both the maximum and 75th percentile are higher for the Top 100. With the same hypothesis as in the previous section the null-hypothesis can not be rejected as this time the p-value 0.4017 is greater than α. The Pearson correlation is again weak with $r = -0.1121$.

To see if there is a relation between the sustainability score and the performance score we use Spearman's Rank Correlation Coefficient [1]. This coefficient is the same that was used by Chan-Jong-Chu et al. to describe the relation between energy consumption and performance score [9]. Applying the test to the Top 100 gives $\rho = 0.12$, but with $p = 0.25$, meaning that the result is not significant. The same applies for the Bottom 100, with $\rho = -0.024$ and $p = 0.83$. We can therefore not find a relation between the sustainability score and the performance score.

5.3.3 Individual cases

University of Birmingham There is one university that achieves the perfect score in the sustainability category: the University of Birmingham, ranked 35 in the Top 100. Upon manual inspection of the website it seems that compression and minification of code is indeed happening. Images are using a modern file format too, but are not responsive. Fonts are using a modern format too, but the used fonts are not web safe. It appears that the way the website is loaded is confusing some of the audits. The index file does not contain the contents of the web page. Instead, the content is loaded through several external files. The received score does not reflect the sustainability of the website well and should have been lower.

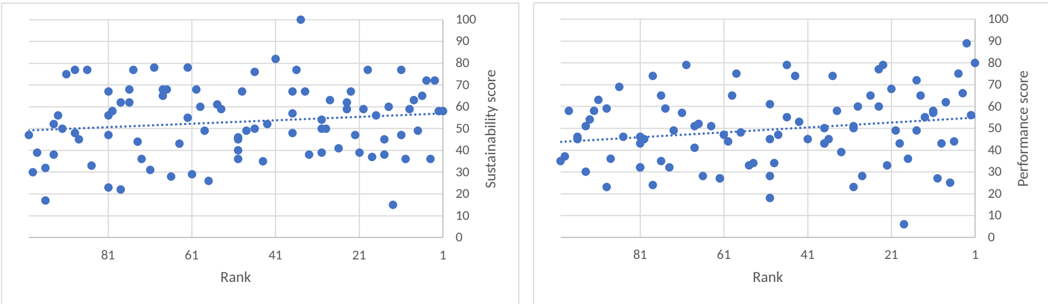

(a) Sustainability against rank (b) Performance against rank

Figure 5.3: Distribution of the performance and sustainability scores of the two groups of universities.

TU Delft In chapter 1 the TU Delft was used as an example of a university with great ambitions for sustainability. This matches with its place in the QS ranking where TU Delft is placed 14th, which is at the top end. This is reflected in the results: the sustainability score is above average at 60, as is the performance score at 65. If we sort the universities in the Top 100 based on their sustainability score the TU Delft is ranked 32nd out of 92. This means that 65% of the universities in the Top 100 perform worse.

Universiti Malaysia Sabah In 2017 Shan et al. wrote that their university had a lot of performance benefits to gain by improving image optimisation, minification and compression [36]. By applying these recommendations the developers of the website could have easily improved the sustainability as well. Based on the results, this has not happened, even though the website has undergone a redesign. The university, found in the Bottom 100, scores 54 for both sustainability and performance.

5.4 Discussion

To measure the relation between the performance score and the sustainability score, we used the Spearman's Rank Correlation Coefficient. We did this to be able to compare our outcome with Chan-Jong-Chu et al., who investigated the relation between the performance score and energy consumption [9]. We had expected to see a similar result. However, whereas Chan-Jong-Chu et al. found a significant correlation, we did not. It is possible that the chosen weights that make up the sustainability score are not well optimised. Another possibility is that the used Lighthouse versions are the cause. For our experiment we used version 11.4.0, the most recent version at the time of development. Chan-Jong-Chu et al. mention that they had used version 3.0, which was released in 2018. But to know for sure, we would have to measure the correlation between the sustainability score and energy consumption.
The chosen audits for the sustainability plugin were based on the Web Sustainability Guidelines and not the patterns listed in chapter 3. The reason for this is that the catalogue has no success criteria, which makes it harder to determine whether a guideline is successfully implemented. To our best knowledge, this study is the first to use the Web Sustainability

Guidelines and to investigate how its guidelines can be implemented in audits. In the future, we expect that more research will focus on how sustainability audits can be designed, but as for now, we consider this a first step.
The TU Delft scored above average on both sustainability and performance, despite the fact that Website Carbon Calculator indicates that its website is dirtier than 61% of all other pages on the internet. It could potentially mean that university websites in general are relatively unsustainable, but more research is needed before this conclusion can be drawn.
The website for the University of Birmingham showed that newly developed audits are not robust. For future experiments the audits should be redesigned to better support websites that use external files to load content. A consequence of the current implementation is that in some cases the sustainability score should be lower. Given that this would be the case for both the Top 100 and Bottom 100 we think that this does not majorly impact the outcome of this experiment.
The experiment in its current form only tells us how sustainable or performant a website is, but not if this was a conscious choice of the university. The universities that top the ranking are large institutes with large budgets. Having a nice website is a way to attract new students and establish a brand. It makes sense that these universities are willing to invest in their website and are thus able to hire good web developers. Whether the universities at the bottom of the ranking are able to make the same kind of investments is questionable, as they are often smaller and located in less developed countries.

5.5 Conclusion

In this chapter we developed a Google Lighthouse plugin that incorporates several guidelines from the Web Sustainability Guidelines. We used this plugin to analyse the websites of universities, split in two groups based on their sustainability, and answer research question **RQ3: Is there a relation between the sustainability of a university and the sustainability of their website?** Between the groups we found a significant difference in sustainability score, but within the Top 100, we only saw a very weak correlation. The difference in performance score between the groups was not significant and the correlation within the Top 100 was also very weak. Furthermore, we were not able to show a relation between the performance score and the sustainability score.

Chapter 6

Discussion

In this thesis sustainable web design is investigated from three different perspectives. In chapter 3, we look through an academic perspective, in chapter 4 through a development perspective and in chapter 5 through an end user's perspective. All perspectives provided new insights, but also confirmed some existing observations. In this chapter we will discuss how a more sustainable web can be achieved based on the insights.

6.1 Academic Perspective

With internet usage growing rapidly worldwide and the need to make it more sustainable, a similar growth in research attention is expected. Nonetheless, the number of publications on this topic between 2012 and 2022 has not grown significantly [15]. It seems that new technologies are a common reason to investigate energy efficiency. OLED-displays led to research on dark themes, new image and video standards led to comparisons to the older standards, and the introduction of WebAssembly gave reason for a comparison with JavaScript. These technical innovations are an important way to achieve a more sustainable web and are easy to understand for developers. Yet, research should also focus on optimising existing technologies as these are still used a lot, as is the case for HTML, CSS and JavaScript.
Media consumption has long been one of the main reasons to use the internet and explains why there is more research on images and videos. But as the internet is used for compute-heavy tasks and as replacement for native applications more and more, it should be viewed as such. Rani et al. does this by taking mobile energy patterns and porting them to web [35]. In this work a first attempt is made at evaluating the patterns through an automated pipeline. One of the outcomes of the study is that one of the patterns, Dynamic Retry Delay, does not work well for web and it illustrates the need to evaluate all patterns.

6.2 Development Perspective

With a lot of knowledge being available, it is up to developers to bring it to practice. In chapter 4 we analysed GitHub repositories to find out how this is done. But almost half of the internet is built through Content Magement Systems (CMS), of which WordPress is the most used [23]. With CMS most of the code is abstracted away so that users with less technological knowledge can also build websites. It is not uncommon for restaurants, barbers or shops to use template websites to which they add their own information. If even web developers are mostly unaware of sustainable web design, it is unlikely that users of a CMS have ever heard of it. The responsibility lays with the CMS in this case to adopt green practices. The Web Almanac has dedicated a chapter to CMS and highlights some differences between available platforms [23, chapter 18]. WordPress's adoption of WebP for images is very limited, whereas another platform, Wix, uses it in roughly 75% of the cases. At the same time, HTML and JS files are over three times as large on Wix compared to WordPress. Implementing green patterns on CMS platforms is the low-hanging fruit of internet sustainability. If the CMS platforms can change the defaults on image formats for example, it can have an overall positive impact.

6.2.1 Tools

The availability of good tools helps developers make good websites. Guidelines and catalogues are part of this and they are even more powerful when integrated in the development environment. The sustainability plugin developed in chapter 5 is an example of how such a tool can make it easier to see where change is needed and what is working well already. Another example is Greenhouse, a Lighthouse plugin that checks if a domain is running on renewable power [20]. By offering such tools developers don't have to know exactly how everything works – which is complicated – but, instead, the tool will indicate what is wrong and how it can be corrected.

6.2.2 Legislation

Governments can play a role in web sustainability too, by enforcing it through laws. Currently, for web accessibility this is already the case in many countries [47]. Developers have indicated that such laws do actually matter in the design process as they have to be followed [25]. Doing the same for sustainable web design can help create the awareness and incentive that is currently absent.

6.3 End User's Perspective

The end user of a website is not powerless, even as the web is egalitarian. There are many ways in which the end user can reduce its footprint while online. From chapter 5 we learned that a sustainable reputation does not necessarily lead to a sustainable website. With a tool like an online carbon calculator it is possible to do these kind of checks for any website. CarbonViz[1] is a browser extension that tracks your browsing behaviour and then gives you details on your online footprint. Such tools can create awareness on how unsustainable using the internet is. Once there is enough awareness among end users, the incentive is created for developers to make their websites more sustainable. Search engines like Google could

[1] https://chromewebstore.google.com/detail/carbonviz/lieikallcilhkihohkajnhmhkiipaaaf

also incorporate sustainability as part of their ranking algorithm, or display a sustainability label next to each search result. That way, it is easier for users of a search engine to make more sustainable choice. For developers it would become part of search engine optimisation, which optimises a website to get a high rank in the search results.

6.4 Conclusion

There are many ways to reach a future in which the internet is more sustainable than in its current form. Researchers can contribute by continuing research on green patterns and other methods to reduce energy consumption for the web. Developers would benefit from the availability of good development tools that make the process of creating sustainable sites easier. Content management systems have the chance to do it by making green patterns the default on their platforms. Yet, all solutions require an awareness that is currently missing. If users of the internet are aware of their emissions, they can adjust their online behaviour to which developers will react. This will create the incentive needed for developers to adopt sustainable web design. And at last, there is the route of policy making, by requiring web developers to comply with sustainable web design legislation.

Chapter 7

Conclusion

Few technologies have become as unmissable and omnipresent as the internet is today. The core technologies behind it were introduced at the end of the 20th century and have been updated through the years to keep up with ever more demanding applications. This has led to an internet that is powerful enough to replace native software applications, allowing people to fully work online. Besides remote working, the internet is used for entertainment, social networking, and as a source of information. But as the benefits are unmistakable, the cost of running a globally used system can not be underestimated. The internet is a large emitter of greenhouse gasses due to all the devices, servers and infrastructure that are involved in keeping it running. With the expected growth of people and devices with access to internet emissions will only increase unless action is undertaken.

In this thesis we investigated what role sustainability currently plays in the process of web design. To this end we formulated three research questions that all investigate a different aspect of the process and together provide a good overview of the current state of sustainable web design and how to improve it.

Research question **RQ1: What are green web design patterns supported by literature?** focused on the academic world and their contributions to web sustainability. In chapter 3 we listed nine energy patterns that are specifically for web and empirically tested to have a positive effect on the sustainability of a website. This catalogue of green patterns is a first attempt at uniting all the work that has been done by other researchers and presenting them in a convenient way for both researchers and developers. The catalogue is a work in progress, as there are many patterns left to be tested before they can be included.

Research question **RQ2: Which design patterns do web developers adopt to improve energy efficiency?** from chapter 4 was targeted at analysing the design decision that developers make during development. To answer the research question we designed a repository mining experiment on GitHub and analysed the discourse in commits, pull requests and issues that contained the words *energy, battery* or *power*. Of the analysed data less than 1% mentioned any form of energy optimisation. Compared to similar studies for different platforms this number is very low [32, 5, 12]. We assume that the low number is a result of the limited awareness and priority on sustainable web design among developers.

The third research question **RQ3: Is there a relation between the sustainability of a university and the sustainability of their website?** looked at the final product, the perspective of the end user. In chapter 5 we measured if there is any relation between the green reputation and the sustainability of a website for 200 universities. For this experiment we developed a Google Lighthouse sustainability plugin that implements nine audits based on the Web Sustainability Guidelines [13]. The results showed that there was a significant difference between the group of universities with the best reputation and the worst, but we were not able to find a correlation between sustainability and performance.

Based on our results we conclude that there is a lot of room for improvement in the field of sustainable web design, mainly due to a lack of awareness and resources. In academia the number of publications on this topic did not increase significantly over the last years. We hope that this will change as there are still many green patterns left to explore. For developers the availability of tools like guidelines, catalogues, plugins and profilers can make it easier to develop sustainable websites. Content Management Systems play a large role in creating a greener internet too, as roughly half the internet is powered by such a system. By creating legislation governments can accelerate the transition to a more sustainable web. Lastly, users of the internet can demand better websites, if they are aware of the consequences of using it. As this awareness is currently absent, other parts in the internet chain like browsers or search engines can help create it by incorporating labels that indicate the sustainability of a web page.

The goal of making a system that is used worldwide but owned by no one more sustainable is not something that can be done overnight. It requires dedication and perseverance and a result will hopefully present itself in the long term.

With this thesis we have contributed:

- a catalogue of green web patterns;
- an analysis of the limited adoption of green patterns in web design;
- an initial implementation of the Web Sustainability Guidelines in the form of a Google Lighthouse plugin;
- an analysis of the relation between reputation and sustainability for 200 university websites;
- all code and data that are publicly available through a replication package.

Bibliography

[1] *Spearman Rank Correlation Coefficient*, pages 502–505. Springer New York, New York, NY, 2008. ISBN 978-0-387-32833-1. doi: 10.1007/978-0-387-32833-1_379. URL https://doi.org/10.1007/978-0-387-32833-1_379.

[2] Victor Agababov, Michael Buettner, Victor Chudnovsky, Mark Cogan, Ben Greenstein, Shane McDaniel, Michael Piatek, Colin Scott, Matt Welsh, and Bolian Yin. Flywheel: Google's data compression proxy for the mobile web. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 367–380, 2015.

[3] Abdurhman Ali Albasir and Kshirasagar Naik. Smow: An energy-bandwidth aware web browsing technique for smartphones. *IEEE Access*, 2:1427–1441, 2014. doi: 10.1109/ACCESS.2014.2365091.

[4] Anders Andrae. Hypotheses for primary energy use, electricity use and co2 emissions of global computing and its shares of the total between 2020 and 2030. *WSEAS Transactions on Power Systems*, 15:50–59, 03 2020. doi: 10.37394/232016.2020.15.6.

[5] Lingfeng Bao, David Lo, Xin Xia, Xinyu Wang, and Cong Tian. How android app developers manage power consumption? an empirical study by mining power management commits. In *Proceedings of the 13th International Conference on Mining Software Repositories*, pages 37–48, 2016.

[6] Barbara Rita Barricelli, Ines Di Loreto, Michele Sciarabba, and Samir Zamouma. Are italian and french public university websites sustainable? In *IFIP Conference on Human-Computer Interaction*, pages 413–417. Springer, 2023.

[7] Tim Berners-Lee. Information management: A proposal. https://www.w3.org/History/1989/proposal.html, 1989.

[8] Yi Cao, Javad Nejati, Muhammad Wajahat, Aruna Balasubramanian, and Anshul Gandhi. Deconstructing the energy consumption of the mobile page load. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(1), 6 2017. doi: 10.1145/3084443. URL https://doi.org/10.1145/3084443.

[9] Kwame Chan-Jong-Chu, Tanjina Islam, Miguel Morales Exposito, Sanjay Sheombar, Christian Valladares, Olivier Philippot, Eoin Martino Grua, and Ivano Malavolta. Investigating the correlation between performance scores and energy consumption of mobile web apps. In *Proceedings of the 24th International Conference on Evaluation and Assessment in Software Engineering*, pages 190–199, 2020.

[10] Thibaud Colas. https://gist.github.com/thibaudcolas/589de9192a919e41b67eefc98e0f4dac#file-marvil07-adapt-thibaudcolas-comparison-table-md. Accessed December 2023.

[11] Luis Cruz. *Tools and Techniques for Energy-Efficient Mobile Application Development*. PhD thesis, University of Porto, 2019.

[12] Luis Cruz and Rui Abreu. Catalog of energy patterns for mobile applications. *Empirical Software Engineering*, 24:2209–2235, 2019.

[13] Alexander Dawson and Tim Frick. Web sustainability guidelines v1.0. https://w3c.github.io/sustyweb/, 2024.

[14] Joao De Macedo, Rui Abreu, Rui Pereira, and Joao Saraiva. Webassembly versus javascript: Energy and runtime performance. In *2022 International Conference on ICT for Sustainability (ICT4S)*, pages 24–34. IEEE, 2022.

[15] Benedikt Dornauer and Michael Felderer. Energy-saving strategies for mobile web apps and their measurement: Results from a decade of research (preprint). *arXiv preprint arXiv:2304.01646*, 2023.

[16] Benedikt Dornauer and Michael Felderer. Web image formats: Assessment of their real-world-usage and performance across popular web browsers. *arXiv preprint arXiv:2310.00788*, 2023.

[17] Benedikt Dornauer, Wolfgang Vigl, and Michael Felderer. On the role of font formats in building efficient web applications, 2023.

[18] Kaushik Dutta and Debra Vandermeer. Caching to reduce mobile app energy consumption. *ACM Trans. Web*, 12(1), 9 2017. ISSN 1559-1131. doi: 10.1145/3125778. URL https://doi.org/10.1145/3125778.

[19] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. The lockdown effect: Implications of the covid-19 pandemic on internet traffic. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 1–18, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450381383. doi: 10.1145/3419394.3423658. URL https://doi.org/10.1145/3419394.3423658.

[20] Green Web Foundation. CO2.JS. https://www.thegreenwebfoundation.org/co2-js/. Accessed April 2024.

[21] Tom Greenwood. *Sustainable Web Design*. A Book Apart, 2021.

[22] Sivert Gullberg Hansen, Henrik Landgraff Granum, Magnus Moen, Ola Hulleberg, Carlos Vicient Monllaó, and Yavuz Inal. Current challenges and barriers in sustainable web design: A qualitative study. In *Proceedings of Mensch und Computer 2023*, pages 375–379. 2023.

[23] HTTP Archive. 2022 web almanac, 2023. URL `https://almanac.httparchive.org/en/2022/`.

[24] Jinyan Hu, Shaojing Song, and Yumei Gong. Comparative performance analysis of web image compression. In *2017 10th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 1–5. IEEE, 2017.

[25] Ola Hulleberg, Henrik Landgraff Granum, Sivert Gullberg Hansen, Magnus Moen, Carlos Vicient Monllaó, and Yavuz Inal. The awareness and practices of web developers toward sustainable web design. In *International Conference on Human-Computer Interaction*, pages 134–145. Springer, 2023.

[26] René Itten, Roland Hischier, Anders SG Andrae, Jan CT Bieser, Livia Cabernard, Annemarie Falke, Hugues Ferreboeuf, Lorenz M Hilty, Regula L Keller, Etienne Lees-Perasso, et al. Digital transformation—life cycle assessment of digital services, multifunctional devices and cloud computing. *The International Journal of Life Cycle Assessment*, 25:2093–2098, 2020.

[27] Samuel Khayo, Dr Ondulo, and Kelvin Omieno. Investigating optimization strategies that determine performance of cms websites in organizations. *IJARCCE*, 7:97–102, 10 2018. doi: 10.17148/IJARCCE.2018.71022.

[28] Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von Werra, and Harm de Vries. The stack: 3 tb of permissively licensed source code. *Preprint*, 2022.

[29] Ding Li, Angelica Huyen Tran, and William GJ Halfond. Making web applications more energy efficient for oled smartphones. In *Proceedings of the 36th International Conference on Software Engineering*, pages 527–538, 2014.

[30] Ding Li, Angelica Huyen Tran, and William GJ Halfond. Nyx: A display energy optimizer for mobile web apps. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, pages 958–961, 2015.

[31] Jennifer Mankoff, Robin Kravets, and Eli Blevis. Some computer science issues in creating a sustainable world. *Computer*, 41(8):102–105, 2008.

[32] Irineu Moura, Gustavo Pinto, Felipe Ebert, and Fernando Castor. Mining energy-aware commits. In *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*, pages 56–67. IEEE, 2015.

[33] Renee Obringer, Benjamin Rachunok, Debora Maia-Silva, Maryam Arbabzadeh, Roshanak Nateghi, and Kaveh Madani. The overlooked environmental footprint of increasing internet use. *Resources, Conservation and Recycling*, 167:105389, 2021. ISSN 0921-3449. doi: https://doi.org/10.1016/j.resconrec.2020.105389. URL https://www.sciencedirect.com/science/article/pii/S0921344920307072.

[34] Quacquarelli Symonds. Qs world university rankings: Sustainability 2024. https://www.topuniversities.com/sustainability-rankings, 2024.

[35] Pooja Rani, Jonas Zellweger, Veronika Kousadianos, Luis Cruz, Timo Kehrer, and Alberto Bacchelli. Energy patterns for web: An exploratory study. *arXiv preprint arXiv:2401.06482*, 2024.

[36] Teh Shan Shan, Joe Henry Obit, Rayner Alfred, and Asni Tahir. Enhancing the performance of university's website for mobile devices based on responsive web design approach. *Advanced Science Letters*, 23(11):10969–10973, 2017.

[37] Steve Souders. High-performance web sites. *Communications of the ACM*, 51(12): 36–41, 2008.

[38] Davide Spadini, Maurício Aniche, and Alberto Bacchelli. Pydriller: Python framework for mining software repositories. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2018, page 908–911, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355735. doi: 10.1145/3236024.3264598. URL https://doi.org/10.1145/3236024.3264598.

[39] Narendran Thiagarajan, Gaurav Aggarwal, Angela Nicoara, Dan Boneh, and Jatinder Pal Singh. Who killed my battery? analyzing mobile browser energy consumption. In *Proceedings of the 21st international conference on World Wide Web*, pages 41–50, 2012.

[40] Estrella Trincado, Antonio Sánchez-Bayón, and José María Vindel. The european union green deal: clean energy wellbeing opportunities and the risk of the jevons paradox. *Energies*, 14(14):4148, 2021.

[41] TU Delft. Sustainable tu delft. https://www.tudelft.nl/en/sustainability. Accessed April 2024.

[42] Mikko Uitto. Energy consumption evaluation of h.264 and hevc video encoders in high-resolution live streaming. In *2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 1–7. IEEE, 2016.

[43] Hüseyin Ünlü and Yeliz Yesilada. Transcoding web pages via stylesheets and scripts for saving energy on the client. *Software: Practice and Experience*, 52(4):984–1003, 2022.

[44] Ahmed Uzair, Mirza O Beg, Hasan Mujtaba, and Hammad Majeed. Weec: Web energy efficient computing: A machine learning approach. *Sustainable Computing: Informatics and Systems*, 22:230–243, 2019.

[45] Max Van Hasselt, Kevin Huijzendveld, Nienke Noort, Sasja De Ruijter, Tanjina Islam, and Ivano Malavolta. Comparing the energy efficiency of webassembly and javascript in web applications on android mobile devices. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering*, pages 140–149, 2022.

[46] Matteo Varvello and Benjamin Livshits. Shedding (some) light on mobile browsers energy consumption. In *TMA*, 2021.

[47] W3C. Web accessibility laws & policies. `https://www.w3.org/WAI/policies/`. Accessed April 2024.

[48] Yuhao Zhu and Vijay Janapa Reddi. Greenweb: Language extensions for energy-efficient mobile web computing. *ACM SIGPLAN Notices*, 51(6):145–160, 2016.