Bachelor thesis
Computing Science

Radboud University

The landscape of (un)balanced choices in cookie consent dialogues in Europe

Author:
Gijs Kopmeiners
s1026838

First supervisor/assessor:
dr. M.G.C. Acar
g.acar@cs.ru.nl

Second assessor:
dr. H.K. Schraffenberger
hanna.schraffenberger@ru.nl

July 17, 2023

Abstract

Website publishers often present cookie consent dialogues to users in order to receive the consent required to place cookies on a user's machine. However, dark patterns, deceptive user interface design techniques, are present within these dialogs, nudging users in to giving their consent. In this study, we look at one instance of dark patterns: not having an option to opt-out on the same layer as an option to opt-in. The objective is to investigate the prevalence and characteristics of this dark pattern in five different European countries. For this, we developed a crawler using the Selenium package in Python, using accept and reject word lists to match the text of web elements. Crawling 23.303 websites, we found 13.522 consent dialogs have the accept option, while only 6.016 have the reject option on the first layer. This indicates that over half of the websites crawled did not have a reject button on the same layer as an accept button, indicating a dark pattern and possibly violating the GDPR.

Contents

Chapter 1

Introduction

Internet users worldwide are flooded with cookie consent dialogues when browsing. Although the primary purpose of cookies is to provide essential functionality such as maintaining a shopping basket between a client and a webshop, they are also used for tracking purposes and collecting user data. These cookies are placed by third parties in websites all across the internet to track your interests and movement trough the web and contain personal information of website users. Website publishers, however, do not always safeguard their users' privacy and data. Following the ePrivacy Directive (ePD) passed in 2009, and later enforced with the General Data Protection Regulation (GDPR) passed in 2016, website publishers have to gain a users' consent in order to place non-essential cookies on a users' machine. [7] The ePD made it mandatory to collect user's consent before any access or storage of non-mandatory data, such as third party cookies. In the case of websites, the consent is usually presented in the form of cookie banners or cookie notices, that pop up when visiting a site and inform the user of data collection. They should provide a meaningful choice on whether to accept or reject such collection. However, even if a website presents a consent and a reject choice, they are often not fair and/or equivalent. These websites use so called dark design patterns to nudge the user into giving consent, even though this might not be in the users interest. [9] The most prevalent pattern currently is obstruction. [21] Obstruction is defined by Gray et al as: "Impending a task flow, making an interaction more difficult than it inherently needs to be with the intent to dissuade an action." In this thesis we look at one case of obstruction: hiding the option to deny consent, especially focusing on cookie banners not having an opt-out or reject button in the same layer as an opt-in or accept button.

We conduct an automated crawl of 23.305 websites in 5 European countries, looking at cookie banners and comparing the prevalence of this instance of obstruction. The crawl is conducted in The Netherlands, Germany, France,

Great Britain and Ireland. These countries are selected as the Netherlands, Germany, France and Ireland are part of the EU and thus fall under EU legislation. Great Britain has recently stepped out of the EU, but still has legislation in place covering the GDPR. While the legality of this instance of obstruction is somewhat a grey area, The Data Protection Authorities (DPA's) of France, the UK and Ireland have explicitly stated they do not allow this form of obstruction. [20] Furthermore, the Cookie Banner Taskforce of the European Data Protection Board (EDPB) has recently stated it does not allow this. [6]

Current attempts at automated auditing crawls often rely on Consent Management Providers (CMPs) to see how cookie banners are implemented to detect dark patterns. [16] [18] However, this method does not work on sites that do not use a CMP and therefore misses the cookie banners on a large number of websites. A better success rate has been achieved with "Priv-accept" [12] or similar methods. [5] Priv-accept uses keywords to look for consent elements that in cookie banners, achieving over 90% success rate. We further improve Priv-accept by extending the consent keyword and creating a reject keyword list to find the reject elements. In chapter 6 we will discuss the implementation as well as the advantages and disadvantages of this method in more detail.

We found that of the 23.207 websites successfully crawled, 13.522 (61%) contained an accept element and 6.016 (28%) contained a reject element on the first layer of the cookie banner. 56% of websites that contained an accept element did not contain a reject element. This means that over half of the websites crawled do not offer a balanced choice, possibly violating the GDPR.

The contributions of this thesis:

1. A concise overview of the current landscape of cookie banners.
2. An extension of Priv-accept [12] to detect reject buttons.
3. Two word lists for every country containing accept and reject phrases encountered in the crawled websites.

In chapter 2, we will discuss the related work related to this paper. In chapter 3, we will explain the methods used in this thesis: How the word lists were created and how the crawler works. In chapter 4, we will present the results gathered by the crawl. In chapter 5, we will discuss the validity of the crawler and the results. In chapter 6, we will give our concluding thoughts.

Chapter 2

Related Work

Websites often place third-party, persistent, marketing cookies that can be used to track users across websites. The website operators often implement a cookie banner to gather consent of users in order to place these cookies. Numerous studies have been conducted to examine the ethics and legality of the design of cookie banners. Researchers have tried to categorize several dark patterns and their impact, as well as ways to automatically research dark patterns and violations of legislation. In this section, we will discuss their findings.

2.1 Dark patterns

Dark patterns are design choices made to implement deceptive functionality that is not in the user's best interest. In 2018, Gray et al. [9] provided an overview of the landscape of dark patterns by defining numerous ways a dark pattern can take shape. They collected examples of dark patterns and did a comparative analysis to define five primary dark patterns: nagging, obstruction, sneaking, interface interference, and forced action. Obstruction is defined as "impeding a task flow, making an interaction more difficult than it inherently needs to be with the intent to dissuade an action". They argued that in the case of obstruction, the user is not able to fully understand all possibilities, which would mean the user can not give informed, freely given consent. Finally, they put responsibility on the designer, arguing that they had an ethical responsibility to make sure they prevent using dark patterns.

Three years later in 2021, Gray et al. published another study discussing dark patterns, through an "interaction criticism" approach [10]. They described the connection between Human-Computer interface (HCI), design, privacy and data protection that is inherent to dark patterns in consent banners. They analyzed designs for their legality and ethical fitness. The main focus was on tracking walls: consent banners that you need to interact with

before the site can be accessed, which are explicitly forbidden by the EDPB. Concluding, they stated the need to engage all aforementioned perspectives to ensure the ethical concerns can be converted to policy.

Utz et al. [22] conducted a study inspecting the cookie banners of a sample of 1000 websites and identified common variables of their user interfaces. They found that 95.8% of the websites either had no option to consent or deny, or only an option to consent in the first layer. The positions, sizes, and "blocking" status (whether the consent notices denies the visitors from accessing the site before a choice is made) were also investigated. Furthermore, they investigated interactions of websites visitors with different versions of cookie banners and found that the aforementioned properties of a cookie banners substantially affects people's behavior. Additionally, they state that many current cookie banners do not offer a meaningful choice to users.

Soe et al. [21] conducted a field study of 300 web-pages. They use the definitions in Gray et al.'s 2018 study [9] to identify numerous categories of dark patterns. This was done manually and each site was examined by two researchers independently. In 43% (129 out of 300) of the cases, there was some form of obstruction. They also found that in 220 of the websites the "deny consent" option is not on the same layer as the "give consent" option. They further defined a dark pattern "Choice cascade", where reaching a deny consent button is only achieved after a number of clicks on buttons like "configure", "learn more", "manage partners", etc.

Nouwens et al. [18] conducted a study researching how Consent Management Platforms (CMPs) designs affect people's consent choices. They scraped the designs of five popular CMP's, yielding 680 CMP implementations. These were then evaluated against European law and regulatory guidance. Among other things, they tested if "Accepting all is as easy as rejecting all". They found only 12.6% of sites had a "reject all" button accessible on the same layer as an "accept all" button. Furthermore, they studied the effect of designs on the consent answer given by users. Conducting an experiment under 40 participants, they found that removing the "reject all" button from the first layer increased the probability of consent by 22 percentage points and conclude that this proves that this practice makes it more likely for users to provide consent, violating the principle of "freely given". Additionally, they concluded that placing controls or information below the first layer renders it effectively ignored.

Graßl et al. [8] conducted two experiments, researching the effects of common design nudges on users' consent decisions and their perception of control. In their first experiment, among other things, they tested the impact of obstruction by presenting the option "Manage options" instead of "Do Not

Agree" during their experiments. Interestingly, they found that people reported more perceived control over their personal data when the "Do Not Agree" option was obstructed by "Manage options". However, it was also stated that because of the generally low levels of perceived control no interpretations could be made about this without more evidence. In their second experiment they tested so called "Bright patterns", essentially reversing the direction of the nudges in dark patterns. They found that these nudges substantially affected people's choices, steering them in the nudged direction. They also mentioned a potential form of conditioning which may present due to because dark patterns being used for longer periods of time, leading users to show behavior not in line with their own interests.

2.2 Legal compliance

While the focus of this paper is not on the legal aspect of the design cookie banners, it is important to introduce some background information to realize the motivation behind this study. Studies discussing the legal aspects of cookie banner design often refer to the GDPR and ePD. The eDP imposes the need for consent for storing and accessing cookies. However, it does not state a form in which this consent should be given.[20] The ePD provides supplementary rules to the GDPR. As placing cookies on a user's device is processing personal data, there must be a legal basis to do so to comply with the GDPR. While the GDPR is a regulation, and therefore directly enforceable in every European country, the ePD is a directive, and is left up to each member state to implement in its own national law.[16] Article 6(1) of the GDPR states 6 legal bases: consent, contract, legal obligations, vital interests of the data subject, public interest and legitimate interest. For any cookies except strictly necessary cookies, websites need consent. Article 7 specifies the conditions for consent: Article 7(3) of the GDPR states that withdrawing consent should be as easy as giving it, and Article 7(4) states that consent should be freely given.

Santos et al. [20] described cookie banners, as a consent mechanism in web applications. They argued that these should be designed and implemented to be compliant with the ePD and GDPR, for which they defined 22 legal requirements. One of these requirements came from their own legal interpretation: "balanced choice" (R13): "From Article 7(4) of the GDPR which states that withdrawing consent should be as easy as giving it, we additionally interpret that the choice between "accept" and "reject" must be consequently balanced (or equitable)." This would mean that the absence of a reject button in the first layer of a cookie banner is not compliant with the GDPR, as the choice between "accept" and "reject" is not balanced. Several DPAs, as well as a recent opinion of the CJEU's Advocate General [3] have

explicitly mentioned this as well. The authors additionally stated that at the time the study was conducted, it was not possible to verify this requirement automatically because of lack of standards in cookie banner design.

Martini et al. [15] also stated that not having a reject button in the first layer is a violation of the GDPR. They argued that as under the GDPR it must be as easy to withdraw as to give consent, the same must a fortiori apply to the initial rejection of approval.

Matte et al. [16] studied IAB Europe's Transparency and Consent Framework (TCF). They did this trough identifying potential legal violations in implementations of cookie banners and detecting such suspected violations by crawling 1.426 websites that contain TCF banners. They detected these violations through analyzing CMPs. While they studied a number of different violations, not having a reject button on the first layer was not one of them.

More recently, the Cookie Banner Taskforce of the EDPB has clearly stated new guidelines on the design of cookie banners.[6] The EDPB has adopted various opinions and guidelines to clarify fundamental provisions of the GDPR and to ensure consistency in the application of the GDPR by DPAs.[20] This report reflects the common denominator agreed by the DPAs in their interpretation of the applicable provisions of the ePrivacy Directive, and of the applicable provisions of the GDPR and reflects a minimum threshold. Here, it is explicitly stated in Article 3 under point 8: "a vast majority of authorities considered that the absence of refuse/reject/not consent options on any layer with a consent button of the cookie consent banner is not in line with the requirements for a valid consent and thus constitutes an infringement"

2.3 Automated audit

While the ethical and legal implications of design choices of cookie banners are starting to concretize, the need for an automated audit process to investigate cookie banners becomes more important. Several studies have been conducted to try and automate this process. In this section we will discuss their methods and compare them to our method.

Utz et al. [22] described several common variables in user interfaces. They manually analyzed 1.000 cookie banners and described the dark patterns they found. However, they did not describe a method to audit cookie banners automatically. In their study, they manually inspected screenshots to see if they contained a consent notice and what the design choices are. Fur-

thermore, they described a method to automatically survey a great number of participants by placing different types of consent notices on a German website. Using this method, they were able to get responses of more than 80.000 unique users.

Matte et al. [16] designed two tools, Cookinspect en Cookie Glasses and ran two automated an semi-automated crawls on 28.257 websites. In this study, they focused on banners that use CMP's that respect IAB Europe's TCF: "a technical specification that allows third-parties to collect and exchange user's consent to data collection and the use of cookies." Cookie Glasses is a tool to enable users to see if consent stored by CMPs corresponds to their choice. The detection of cookie banners was done trough Cookinspect, which detects the presence of a IAB Europe's TCF banner by checking whether a `__cmp()` function is defined on a website. All CMPs that use IAB Europe's TCF must implement such a function. This way, they could run an automated crawl, detecting several possible violations: consent stored before choice, no way to opt out, pre-selected choices and non-respect of choice. These violations could be detected with a high certainty. However, this method had a major drawback: it can only detect banners that use a CMP that respects IAB Europe's TCF. In their study, they found only 1.426 of the 28.257 (5%) crawled websites implemented this, greatly limiting their scope. Our method does not rely on these CMPs and can detect cookie banners on any site.

Nouwens et al. [17] presented Consent-O-Matic, a browser extension that automatically handles cookie banners. It is able to set users' data processing preferences and bypasses the interfaces of existing consent pop-ups. Consent-O-Matic also uses CMPs to detect cookie banners. Afterwards, depending on the CMP found, it executes a set of actions to submit consent equal to the user's preferences.

Soe et al. [21] manually analyzed 300 cookie banners from Scandinavian and English news outlets. In contrast to Matte et al. [16] and Nouwens et al. [18], they focused specifically on features that were hard to detect automatically. They investigated the existence and variety of dark patterns categorized by [9], the possibility for the user to not give consent, the location of the consent notice on the screen and the complexity of the consent notice and found that all websites employ some level of unethical practices. They stipulated the need for clear dark pattern categorization with automatically identifiable characteristics, so an automated audit becomes more feasible.

Hausner et al. [11] tried using machine learning approaches to detect dark patterns as generically as possible. They found that out of 4000 German web sites, around 2800 cookie banners could be extracted and analyzed. However, they do not report anything on the effectiveness or validity of their

approach other than the statement that "the implemented framework is powerful enough to detect cookie banners on a wide range of web pages". To the best of our knowledge, this is still an ongoing project.

Kampanos et al. [13] conducted an empirical study of more than 17,000 websites in Greece and the UK, collecting more than 7,000 cookie banners. They used OpenWPM and an extensive list of CSS selectors cookie banners use from "I don't care about cookies" [1] to identify the cookie banners. Afterwards, they categorized the options within the cookie banners in four categories: Affirmative, Negative, Informational and Managerial. This was done with matching the text of the options with a list of Affirmative, Negative, Informational and Managerial phrases. Among other things, they researched the privacy options in cookie banners. They found that while 95% of Greek and 88% of UK websites had an affirmative option (to opt in), only 20% of Greek and 6% of UK websites had a negative option (to opt out).

Aerts [5] used an extension on the OpenWPM crawler to automatically detect consent and reject elements on websites and record their width, height, and background color attributes. A set of strings, created by manually analyzing a partial list of the top websites of the Netherlands and Belgium, was used to detect these elements. An XPath query is executed to search for elements that contain a string in their list is identified as a consent element. This is very similar to the method used in this paper. However, this method has several drawbacks which we try to address and improve in this paper.

Firstly, their method of collecting strings resulted in a greatly limited set of words, using on average only 10 accept strings and 17 reject strings. Our method extends those lists to 200 to 400 strings. This results in a more accurate detection of consent and reject elements. Secondly, they investigate if a certain string is present in the text of an element, not if it is exactly equal. This increases the possibility of detecting false positives, i.e. the existence of the word "prima" (meaning something like "okay" in Dutch) in the string "Skip to primary content", falsely classified this element as a consent element. Our method only looks at the full text of an element, reducing false positives.

Jha et al. [12] used Priv-accept to detect consent elements in a cookie banner. Priv-accept uses a keyword list to looks for accept elements. They stated that compared to other methods, Priv-Accept is the best approach, finding the correct element (when present) in 90% of cases. That is why in this paper, we used a variation of Priv-accept to detect consent and reject elements. We further extended Priv-accepts consent word list and create a reject word list.

During our study, Rasaii et al. [19] published a study in which they developed BannerClick, a tool to automatically detect, accept, and reject cookie banners. They report an accuracy of of 99% for detecting, 97% for accepting and 87% for rejecting cookie banners. Compared to Priv-accept used by Jha et al. [12], BannerClick is able to detect reject elements. Our implementation of Priv-accept however, is also able to detect reject elements. Where BannerClick first detects the banner and then searches the elements in the banner, Priv-accept searches all elements in the DOM, resulting in detecting more elements identified as accept elements when they are not. This is not addressed in our study. They also mentioned Priv-accept uses only English words and BannerClick works in twelve different languages. Our implementation of Priv-accept works in four different languages.

Chapter 3

Methods

3.1 Dataset

In this research, we use a large set of top websites from several European countries: The Netherlands, Germany, France, The United Kingdom (UK) and Ireland. As stated before, these countries were chosen because The Netherlands, Germany, France and Ireland fall under EU legislation. The UK still has legislation in place covering the GDPR. Additionally, all countries, except France, speak a Germanic language, aiding in translating the word lists (see 4.4 for more information). CNIL, the French DPA, has recently fined Google and Facebook a total of 210 million euros [4] for not having a reject option next to an accept option, making it an interesting country to add to the dataset. We generate a Tranco [14] list for each country. Tranco is a top website list ranking based upon the rankings of Alexa, Umbrella, and Majestic. We use generated Tranco lists with two specific configurations:

1. Only pay-level domains were retained.
2. Only domains included in the Chrome User Experience Report of April 2023, present in the dataset for the corresponding country, were retained.

Only including the pay-level domain ensures that we only get one website per single user or organization, e.g. no "drive.google.com" and "translate.google.com" but only "google.com". This was done because it is likely that the same organisation employs the same kind of cookie banner. Only including domains included in the Chrome User Experience Report of a certain country ensures that we only get domains that are accessed from that country. These Tranco lists can be found in the Appendix. Then, we use the country code top-level domain (ccTLD) as a filter to include only websites of a certain country. The ccTLD's used were:

Country	ccTLD
The Netherlands	nl
Germany	de
France	fr
The UK	co.uk
Ireland	ie

Figure 3.1: ccTLD used

The top 5000 sites of each list were collected and these formed our dataset. This dataset can be found on the github under `websites`. Ireland's top list consisted of only 3.305 domains, making the total dataset consist of 23.305 unique domains.

3.2 System setup

The crawl was run sequentially on one PC. A Virtual Private Network, Mullad VPN [2], was used to simulate the crawl being run from different locations. All crawls were run trough a VPN server in their corresponding country to simulate visiting the site from the corresponding country.

Country	City where VPN server was located
The Netherlands	Amsterdam
Germany	Berlin
France	Paris
The United Kingdom	London
Ireland	Dublin

Figure 3.2: City in which the VPN server was located

3.3 Priv-accept implementation

At the core, our crawler uses an implementation of Priv-accept, a Selenium-based crawler described by Jha et al. [12] This is a relatively lightweight crawler that uses word lists to search accept elements in a website and tries to click on these elements. We extend this implementation with added functionality, like drawing a red border around the elements found, aiding in word list creation and manual verification. Additionally, functionality to look for reject elements and record data of the elements found was also added. We implement this crawler in Python 3, using Selenium as framework and Firefox as the automated browser. The code used in this research has been made

public on: `https://github.com/G1zm0K/priv-accept-reject`. It records for each site:

1. The target URL (string): `url`
2. If it found an accept element (boolean): `accept-found`
3. The type of the accept element (string): `accept-type`
4. The text in the accept element (string): `accept-text`
5. If it found an reject element (boolean): `reject-found`
6. The type of the reject element (string): `reject-type`
7. The text in the reject element (string): `reject-text`
8. If an error occurred while visiting or crawling the site (boolean): `error`

The global steps the crawler takes for each site are as follows:

1. Visit the target URL and wait for 3 seconds for the cookie banner to load.

```
driver.get(site)
sleep(3)
```

2. Try to find an accept element:
   (a) Find all elements in the DOM that are a button, a, p, div, span or form.

```
elements = driver.find_elements(By.CSS_SELECTOR,
↪  selectors)
```

   (b) For each element, match the text of that element to the keywords of the accept word list. The text of the element is converted to lowercase and stripped of whitespaces, exclamation marks, etc. to aid in matching.

```
for e in elements:
    if e.text.lower().strip(" >!\n>") in words_list:
```

   (c) If a match is found, draw a red border around the element and store the type and text of the element.

```
driver.execute_script("arguments[0].style.border='4px
↪  solid red'", e)
return
↪  (True,str(e.tag_name),str(e.text.lower().strip("
↪  >!\n>")))
```

(d) If an initial match is not found, try again for each iframe on the site.

```
if not accept_found:
    iframes = driver.find_elements(By.CSS_SELECTOR,
    ↪  "iframe")
    for iframe in iframes:
        driver.switch_to.frame(iframe)
        (accept_found,accept_type,accept_text) =
        ↪  find_banner(driver,selectors,accept_words_list)
        driver.switch_to.default_content()
```

3. Try to find a reject element, this is the same as finding an accept element, except using the reject word list instead of the accept word list.

```
(reject_found,reject_type,reject_text) =
↪  find_banner(driver,selectors,deny_words_list)
```

4. Sleep for one second to make sure the red borders are drawn and take a screenshot.

```
sleep(1)
driver.save_screenshot('screenshots/' + str(index) +
↪  '.png')
```

5. Store all data in a pandas dataframe

```
new_entry = {'index': index, 'url': site, 'accept-found':
↪  accept_found, 'accept-type': accept_type,
↪  'accept-text': accept_text, 'reject-found':
↪  reject_found, 'reject-type': reject_type,
↪  'reject-text': reject_text, 'error': False}
new_df = pd.DataFrame(new_entry, index=[0])
df = pd.concat([df, new_df], ignore_index=True)
```

3.4 Word list creation

In order for Priv-Accept to work with accept as well as reject elements, two separate word lists were needed (accept and reject words and phrases) for four different languages: Dutch, German, French, and English. The UK and Ireland share lists because both of these countries use English. These separate lists contain the corresponding keywords and phrases. To create these lists, we started with the words and phrases in the word lists used by Aerts. [5] These word lists contained 18 to 29 accept or reject words per list, of which about half are English. To improve these lists, we ran an initial crawl with these word lists on the top 500 websites for each of the five countries in the dataset, marking a red border around the elements found and taking a screenshot. We manually inspected the screenshots and added the text that was in accept and reject elements if it was not already in the lists, essentially scraping the accept and reject elements of 2,500 websites.

Subsequently, for each word list, it was translated into the other three languages. All the word lists in the same language were merged, removing duplicates, to make two extensive word lists for each language. Afterwards, we manually filtered out the words that were not likely to be used in consent notices. Words with a double meaning sometimes got lost in translation. For example, "grant" as in "grant permission" was translated to "beurs" in Dutch, meaning a scholarship. The judgement of this was based on the researchers' knowledge of these languages. The filtered lists contained 214 to 324 word and phrases. These lists can be found on `https://github.com/G1zm0K/priv-accept-reject` under `wordlists`. For the Dutch, French, and German crawls, the English word list was combined with the word list of the country, as part of the consent notices were in English.

Chapter 4

Results

In this chapter we first present the results of the crawl. In chapter 6 we discuss the validity of and interpret these results.

We are interested in the number of accept and reject elements we can find in a cookie banner. If an accept element can be found, but a reject element can not, this would indicate that the accept and reject elements are not on the same layer (if there even is a reject element at all). We are also interested in the words and types of elements used, as this would give more insight in how cookie banners are constructed and how the crawler could be improved.

A total of 23.304 sites were crawled, of which 1.097 returned an exception, either while loading the page or when looking for elements. These sites were not included in any of the rest of the analysis. A total of 13.883 cookie banners were identified. Cookie banners were identified if either an accept or reject element was present on the site. 13.522 websites contained an accept element, and a total of 6.016 websites contained a reject element. Relatively, in France, websites had the most reject elements with 1916 reject elements and 2997 accept elements (63%) In the United Kingdom, websites had relatively the least reject elements, with 818 reject elements and 2915 accept elements (28%).

Country	Total websites	No error	Banners	Accept el.	Reject el.
The Netherlands	5000	4888	3010	2953	1117
Germany	5000	4808	2860	2763	1337
France	5000	4802	3120	2997	1916
The UK	5000	4543	2974	2915	818
Ireland	3305	3166	1919	1894	828

Figure 4.1: Accept and reject elements per country

Figure 4.2: Accept vs reject elements by country

We are also interested in the Tranco rank to make sure our dataset is distributed similarly for each country. Firstly, becauase a country with only very popular websites in its dataset compared to a country with only very non-popular websites in its dataset could distort the results. Secondly, to see if there is a difference in the number of websites with accept or reject elements compared to popularity. In figure 4.3 the rank distributions per country are shown to validate that the distributions are similar for each country.

Country/Tranco Rank	Min	Max	Median	Average
The Netherlands	1.042	135.644	109.314	99.490
Germany	138	120.357	79.297	73.488
France	292	152.741	113.187	102.092
The United Kingdom	102	184.607	126.024	115.512
Ireland	1.939	90.007	82.398	75.573

Figure 4.3: Tranco rank distributions per country

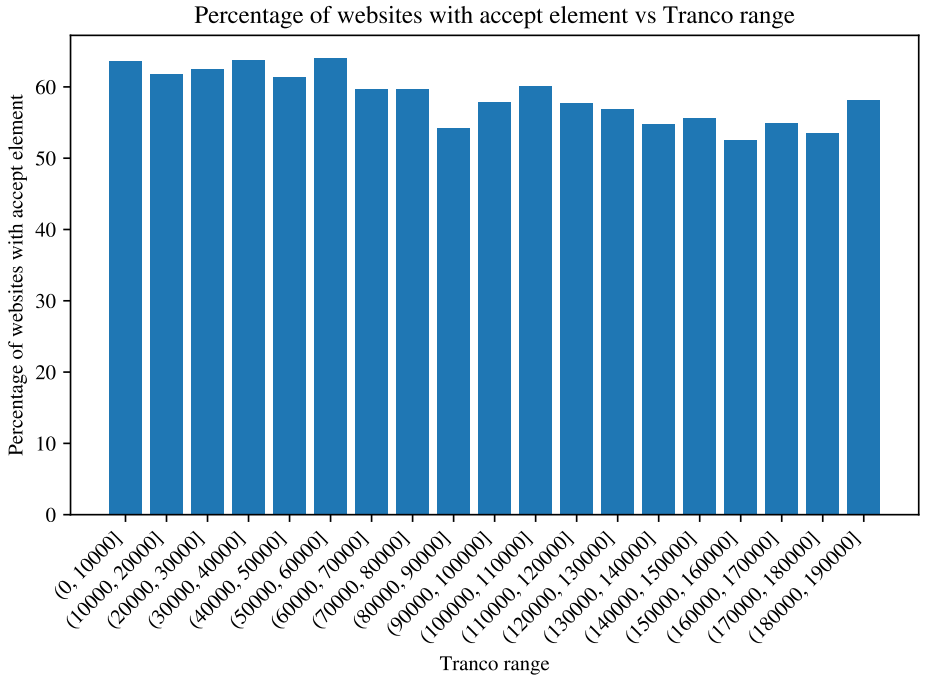

Figure 4.4: Percentage of websites with accept element vs Tranco rank

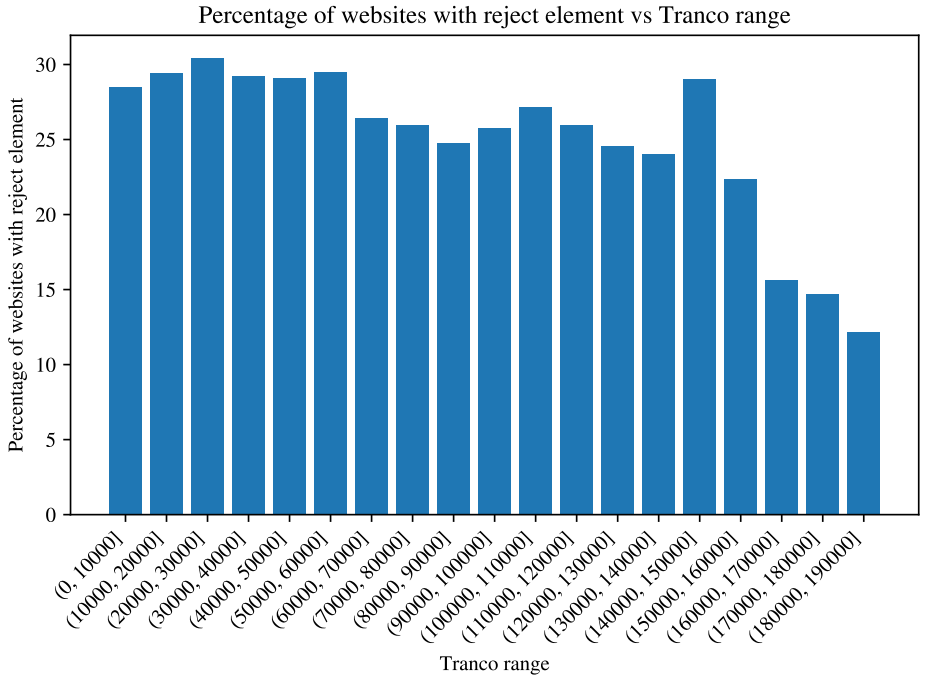

Figure 4.5: Percentage of websites with reject element vs Tranco rank

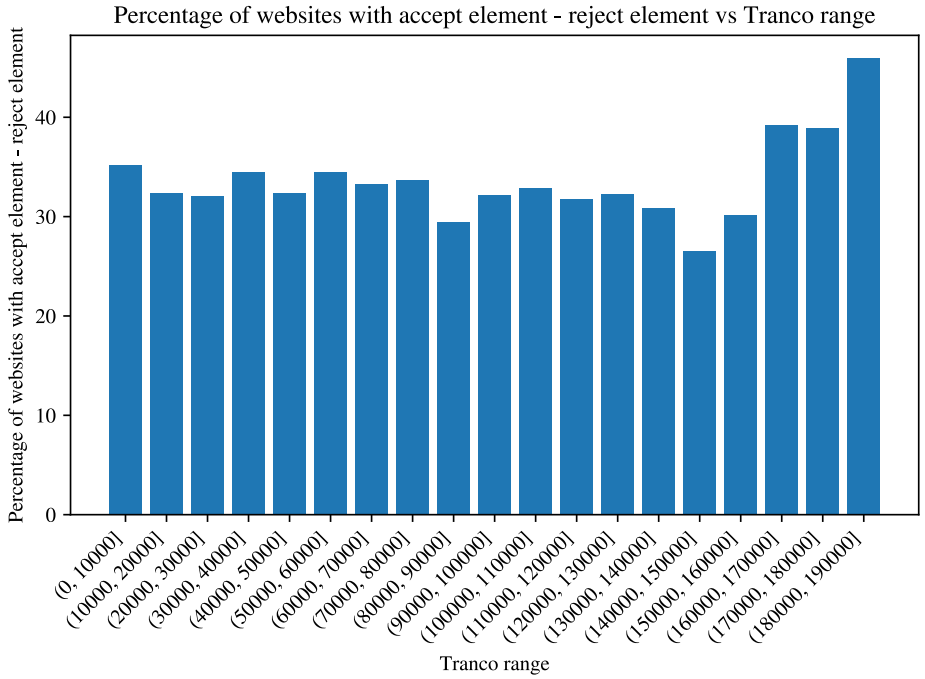

Figure 4.6: Difference between percentage of websites with accept and reject elements vs Tranco rank

The separate graphs for each country can be found in the appendix.

On average, 87.2 unique accept words and phrases and 70.4 unique reject words and phrases were found per country.

Country	No. unique accept words	No. unique reject words
The Netherlands	107	87
Germany	107	86
France	82	70
The UK	83	62
Ireland	57	47

Figure 4.7: Number of unique accept and reject words by country

Below, the top 10 accept and reject words per country are shown. The complete graphs can be found in the appendix.

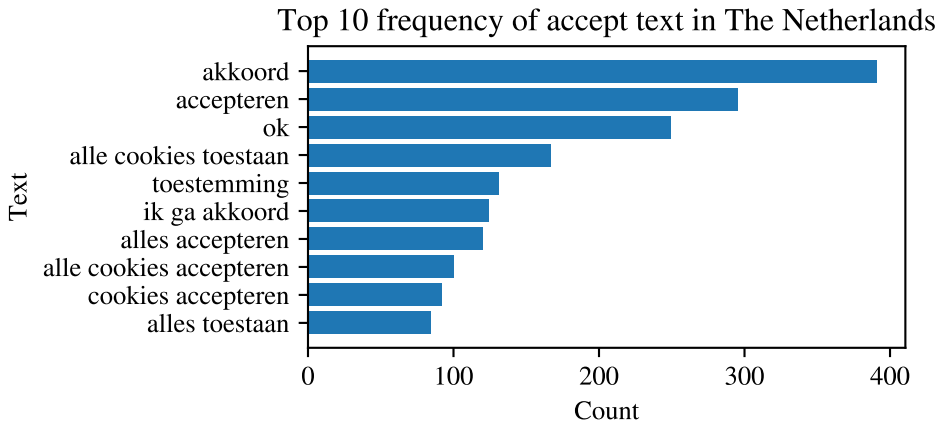

Figure 4.8: Top 10 Frequency of accept text in The Netherlands

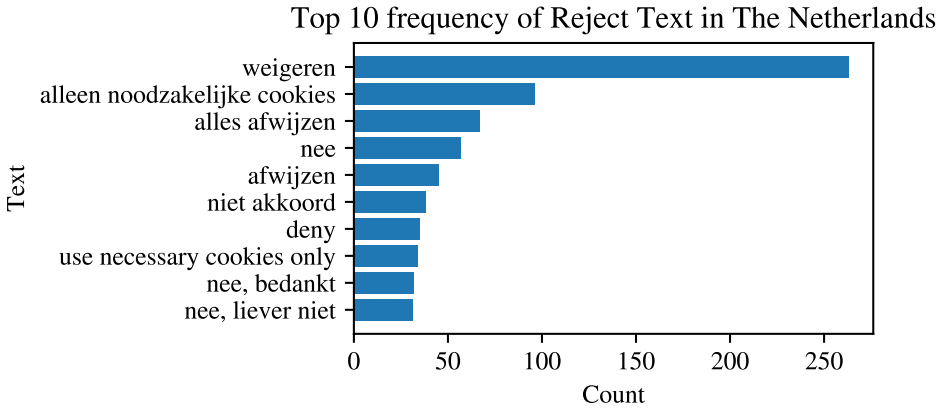

Figure 4.9: Top 10 Frequency of reject text in The Netherlands

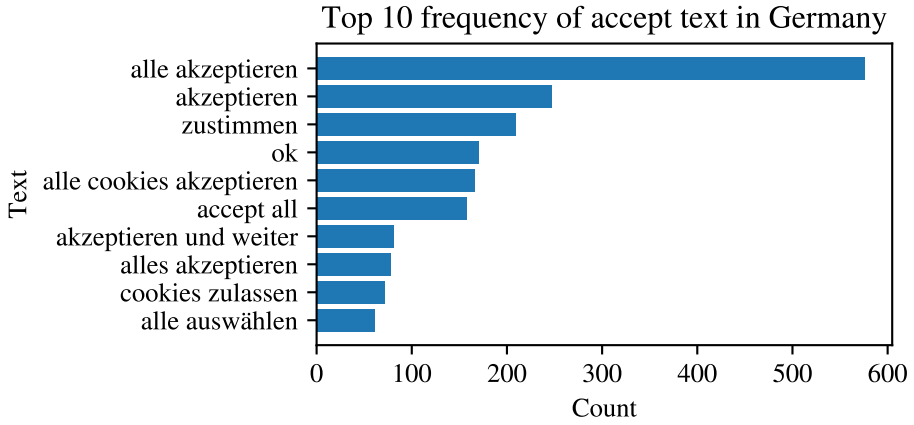

Figure 4.10: Top 10 Frequency of accept text in Germany

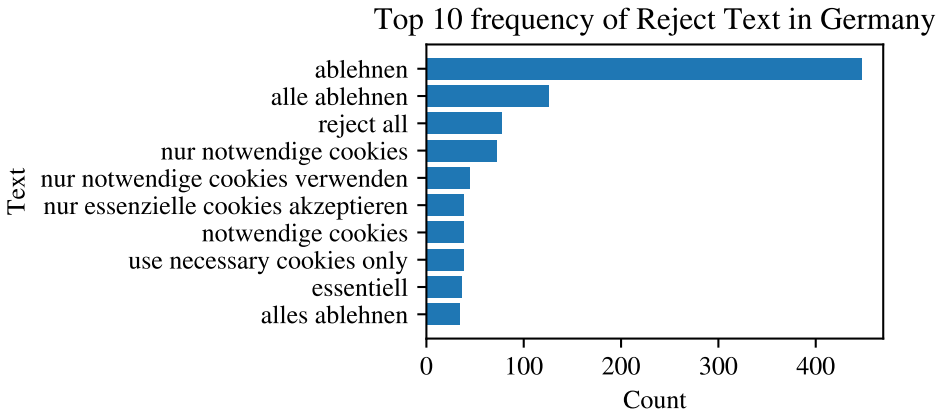

Figure 4.11: Top 10 Frequency of reject text in Germany

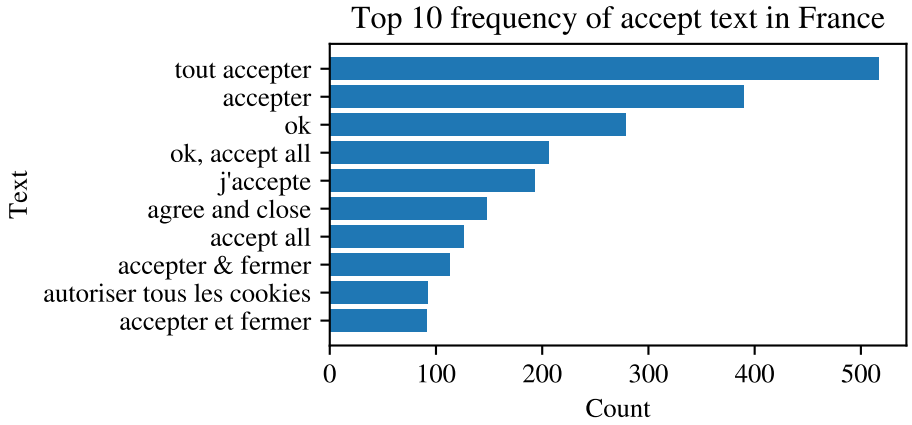

Figure 4.12: Top 10 Frequency of accept text in France

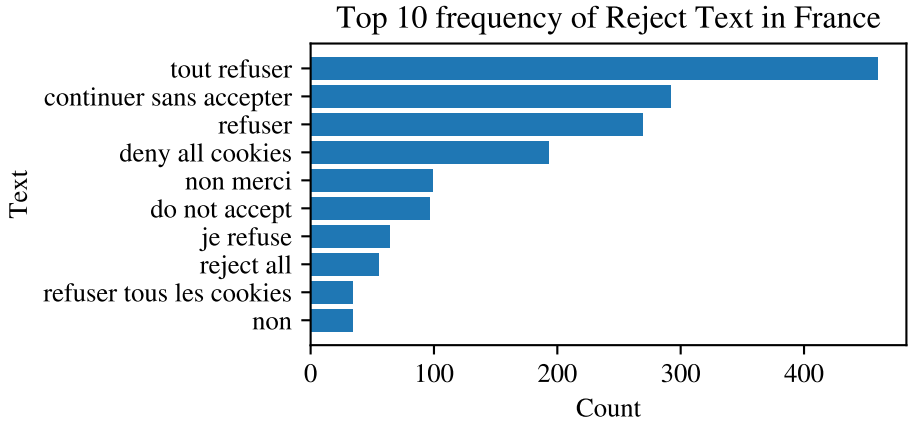

Figure 4.13: Top 10 Frequency of reject text in France

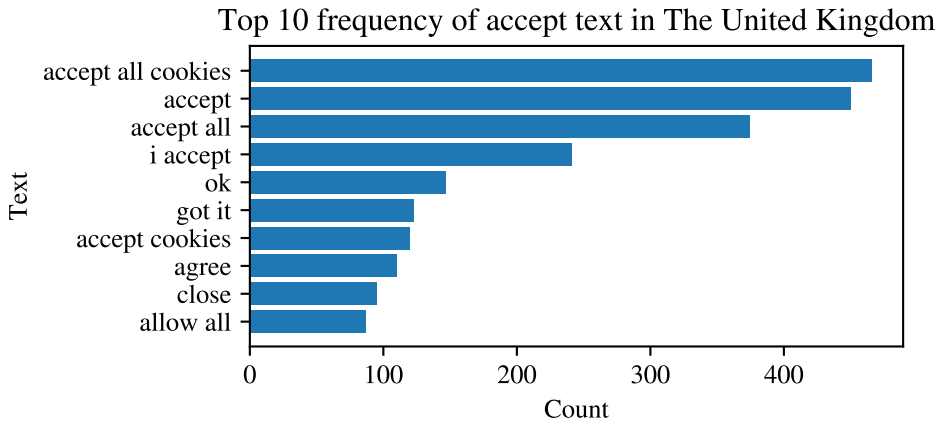

Figure 4.14: Top 10 Frequency of accept text in The United Kingdom

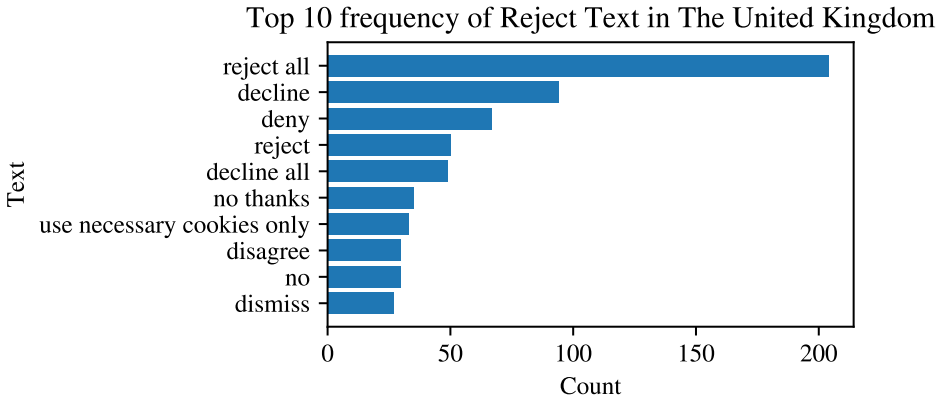

Figure 4.15: Top 10 Frequency of reject text in The United Kingdom

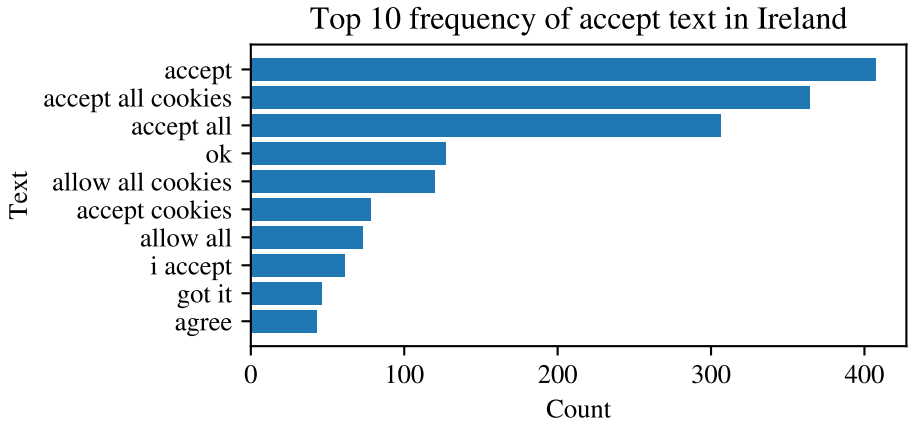

Figure 4.16: Top 10 Frequency of accept text in Ireland

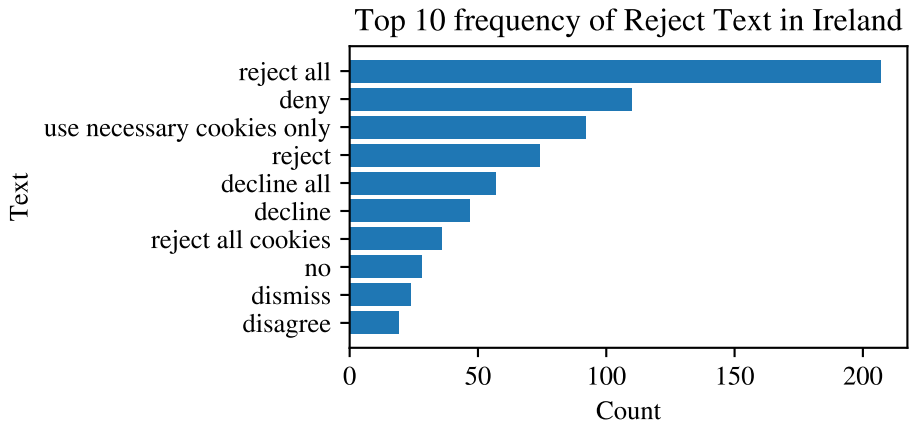

Figure 4.17: Top 10 Frequency of reject text in Ireland

Lastly, we are interested in the type of the accept and reject elements. The distribution per country cam be found in the appendix. The general distribution of types can be seen below:

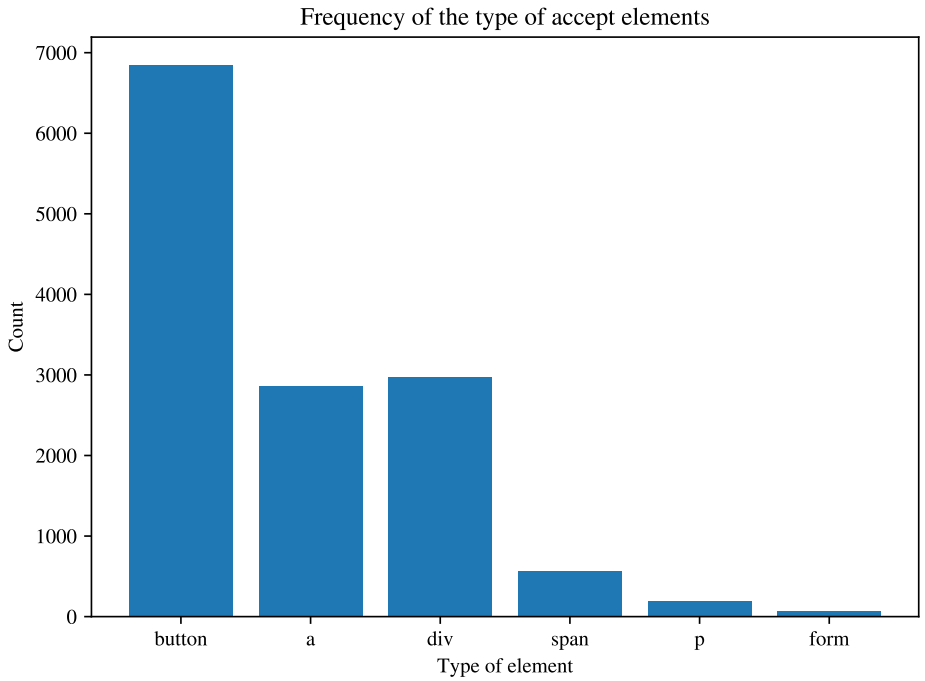

Figure 4.18: Frequency of the type of accept elements

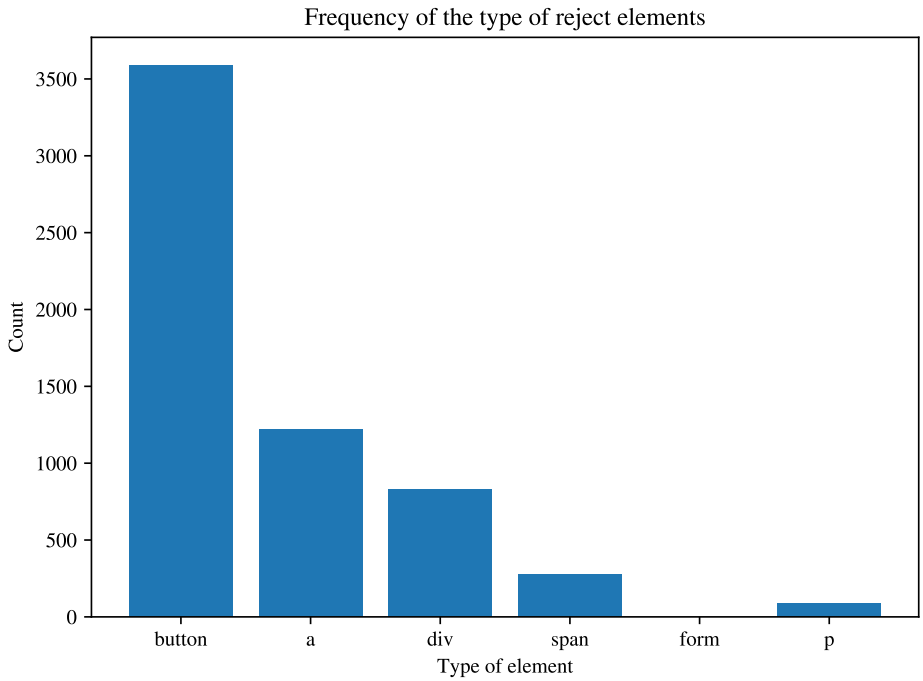

Figure 4.19: Frequency of the type of reject elements

Chapter 5

Discussion

5.1 Validation

To validate the correctness of the accept and reject detection, the screenshots of a random sample of 50 websites for each of the five countries, in total 250 websites, were manually inspected and compared to the collected data. For both accept and reject elements we tested for false positives (FP) and false negatives (FN). FP: if the data said an accept or reject element was found but the screenshots show that element is not an accept or reject element. FN: if an accept or reject element wasn't found while it was present on the site. The results can be seen in figure 5.1 below.

Country	Banners	FN Accept	FP Accept	FN Reject	FP Reject	Total
The Netherlands	34	3	4	2	1	10
Germany	31	8	2	4	1	15
France	32	2	3	3	1	9
The UK	29	3	1	0	0	4
Ireland	32	0	1	0	1	2
Total	158	16	11	9	4	40

Figure 5.1: False positives and negatives of accept and reject elements

These false positives and negatives are higher than initially expected. This could be attributed to the relatively small size of the validation dataset (only 1%) of the total dataset which is potentially not representative enough. However, the anomalies are not very concerning. France doesn't have a big number of false positives with its reject buttons, which could have indicated that the high percentage of websites with accept elements in France would be due to false positives. Germany has a high number of 8 false negatives in the

accept elements. After further inspection, a good explanation couldn't be found, as the elements that weren't found did not show any properties out of the ordinary. However, Germany does not show a large shortage of websites with accept elements, so we argue that this does not heavily influence the results and could only be present in the validation dataset.

5.2 Discussion of the results

In this section, we will discuss the results and their implications. We found significant differences in the number of accept and reject elements per country, as well as significant differences in the relative number of reject elements per accept element per country. As can be seen in figure 4.1, websites in France had the highest percentage of reject elements (41% of websites contained one). This could be because of the fines [4] the French DPA has imposed recently, motivating website publishers to adhere to cookie banner guidelines. Websites in the United Kingdom had the lowest percentage of reject elements (18% of websites contained one). This could be because a bigger portion of the websites in the United Kingdom have low Tranco ranks compared to other countries. Websites on these lower could have a lower percentage of reject elements in any country. However, this hypothesis needs further research as there were too little websites with low ranks in this thesis to validate this. Also, the recent exit of the United Kingdom out of the European Union might be a factor, causing websites publishers to care less about data protection. However, the legislation implemented because of the GDPR is still in place. A great number of websites in the countries crawled do not have a reject button on the same layer as an accept button (74%). This is concerning, as not having a reject button on the same layer as the accept button violates the GDPR, as discussed in section 2.2, as it does not constitute freely given consent.

At first glance, the number of reject elements seems to decrease as the websites get less popular based on figure 4.6. Basing popularity on the Tranco rank (with Tranco rank 1 as the most popular) there is an increase of cookie banners without a reject button in the lower rankings (160.000 and lower), while the number of cookie banners with an accept button stays mostly the same. However, as can be seen in the appendix, ranks lower than 160.000 are only populated by the United Kingdom, which had the lowest number of websites with a reject button. This would indicate that Tranco rank does not have a significant impact on the prevalence of websites with reject elements. However, more research of websites with a lower Tranco rank is needed to confirm or refute this hypothesis. When inspecting the individual graphs more closely, the number of websites with cookie banners decrease in every country except Germany. There does not seem to be a significant increase

of cookie banners without a reject button in the lower rankings.

All distributions of accept and reject texts over all countries follow the same "long-tail" distribution. The first few accept or reject texts are present in the majority of websites. Furthermore, on average, only 87.2 unique accept words and phrases and 70.4 unique reject words and phrases were found per country. Not counting English words in non-English countries, on average 66.6 unique accept words and 51.2 reject words were used. This means only a fraction of words in the lists were used. We provide a list of all the words encountered during the crawl, to aid further research in this area. This list can be found on the GitHub.

The distribution of the frequency of the type of accept and reject elements all look very similar. Buttons are by far the most used. Some accept elements were of the type "form", but no reject elements were of this type. These insights could also aid further research in this area.

Chapter 6

Conclusions

In this thesis, we studied the prevalence of a dark pattern: not having a reject button on the same layer as an accept button. Crawling 23.303 websites in 5 European countries, we constructed a concise overview of the current landscape of cookie banners. We found 61% of websites contained an accept element and only 28% contained a reject element. Furthermore, we found a significant difference in the prevalence of this dark pattern in different countries. A concerning number of websites that do implement an accept button, do not implement a reject button on the same layer. This means users are nudged into giving consent and this is not compliant with the GDPR and the conclusions of the cookie banner taskforce of the EDPB. Furthermore, we did not see a significant difference in the number of accept or reject elements with the decrease of popularity of websites.

Future work could be aimed at extending this crawl to more countries and possibly using our wordlists to enhance either Priv-accept [12] or BannerClick [19] as part of an automated audit for cookie banners. Another direction is a better classification of different cookie banner practices.

Bibliography

[1] I don't care about cookies. https://www.i-dont-care-about-cookies.eu/.

[2] Mullad VPN. https://mullvad.net/en.

[3] Opinion of Advocate General Szpunar delivered on 21 March 2019(1). https://curia.europa.eu/juris/document/document.jsf?docid=212023&doclang=en.

[4] Cookies: the CNIL fines GOOGLE a total of 150 million euros and FACEBOOK 60 million euros for non-compliance with French legislation. https://www.cnil.fr/en/cookies-cnil-fines-google-total-150-million-euros-and-facebook-60-million-euros-non-compliance, 2021.

[5] Koen Aerts. Cookie Dialogs and Their Compliance: The Quest for an Automated Audit Process to Enhance Privacy regulation. https://research.ou.nl/en/studentTheses/cookie-dialogs-and-their-compliance-the-quest-for-an-automated-au, 2021.

[6] EDPB Cookie Banner Task Force. Report of the work undertaken by the Cookie Banner Taskforce. Technical report, 2023.

[7] European Commission. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), 2016.

[8] Paul Graßl, Hanna Schraffenberger, Frederik Zuiderveen Borgesius, and Moniek Buijzen. Dark and Bright Patterns in Cookie Consent Requests. *Journal of Digital Social Research*, 3, 2021.

[9] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. The dark (patterns) side of UX design. In *Conference on Human Factors in Computing Systems - Proceedings*, volume 2018-April. Association for Computing Machinery, 4 2018.

[10] Colin M. Gray, Cristiana Santos, Nataliia Bielova, Michael Toth, and Damian Cliford. Dark paterns and the legal requirements of consent banners: An interaction criticism perspective. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, 5 2021.

[11] Philip Hausner and Michael Gertz. Dark patterns in the interaction with cookie banners. `https://arxiv.org/abs/2103.14956`, 2021.

[12] Nikhil Jha, Martino Trevisan, Luca Vassio, and Marco Mellia. The Internet with Privacy Policies: Measuring the Web Upon Consent. *ACM Transactions on the Web*, 16(3), 9 2022.

[13] Georgios Kampanos and Siamak F. Shahandashti. Accept All: The Landscape of Cookie Banners in Greece and the UK. `https://www-users.york.ac.uk/~sfs521/papers/KS21-Cookie-Banner-UK-Greece-IFIP-SEC-2021.pdf`, 4 2021.

[14] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium*, NDSS 2019, 2 2019.

[15] Mario Martini and Christian Drews. Making Choice Meaningful-tackling Dark Patterns in Cookie and Consent Banners through European Data Privacy Law. `https://ssrn.com/abstract=4257979`, 2022.

[16] Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do Cookie Banners Respect my Choice? Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. `https://arxiv.org/abs/1911.09964`, 11 2019.

[17] Midas Nouwens, Rolf Bagge, Janus Bager Kristensen, and Clemens Nylandsted Klokmose. Consent-O-Matic: Automatically Answering Consent Pop-ups Using Adversarial Interoperability. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, 4 2022.

[18] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, 4 2020.

[19] Ali Rasaii, Shivani Singh, Devashish Gosain, and Oliver Gasser. Exploring the cookieverse: A multi-perspective analysis of web cookies. In *Proceedings of the 2023 Passive and Active Measurement Conference*, March 2023.

[20] Cristiana Santos and Nataliia Bielova. Are cookie banners indeed compliant with the law? Deciphering EU legal requirements on consent and technical means to verify compliance of cookie banners. `https://ssrn.com/abstract=4206054`, 2019.

[21] Than Htut Soe, Oda Elise Nordberg, Frode Guribye, and Marija Slavkovik. Circumvention by design – dark patterns in cookie consents for online news outlets. `http://arxiv.org/abs/2006.13985`, 6 2020.

[22] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un)informed Consent: Studying GDPR consent notices in the field. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 973–990. Association for Computing Machinery, 11 2019.

Appendix A

Appendix

Tranco lists:

IE: `https://tranco-list.eu/list/992J2`
GB: `https://tranco-list.eu/list/Y5JZG`
FR: `https://tranco-list.eu/list/PZP5J`
DE: `https://tranco-list.eu/list/W9GV9`
NL: `https://tranco-list.eu/list/LYZ84`

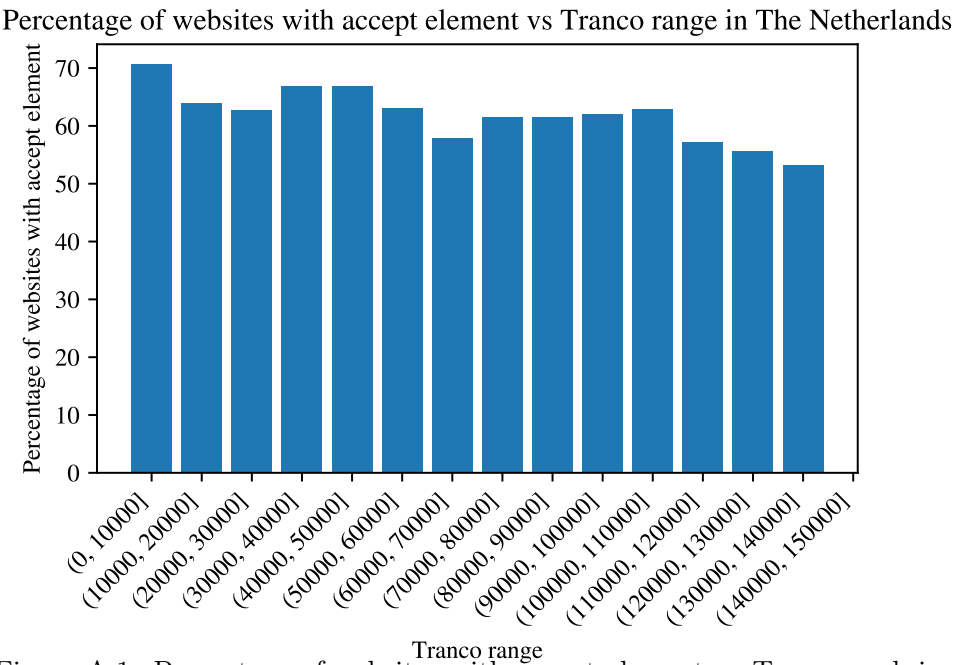

Figure A.1: Percentage of websites with accept element vs Tranco rank in: The Netherlands

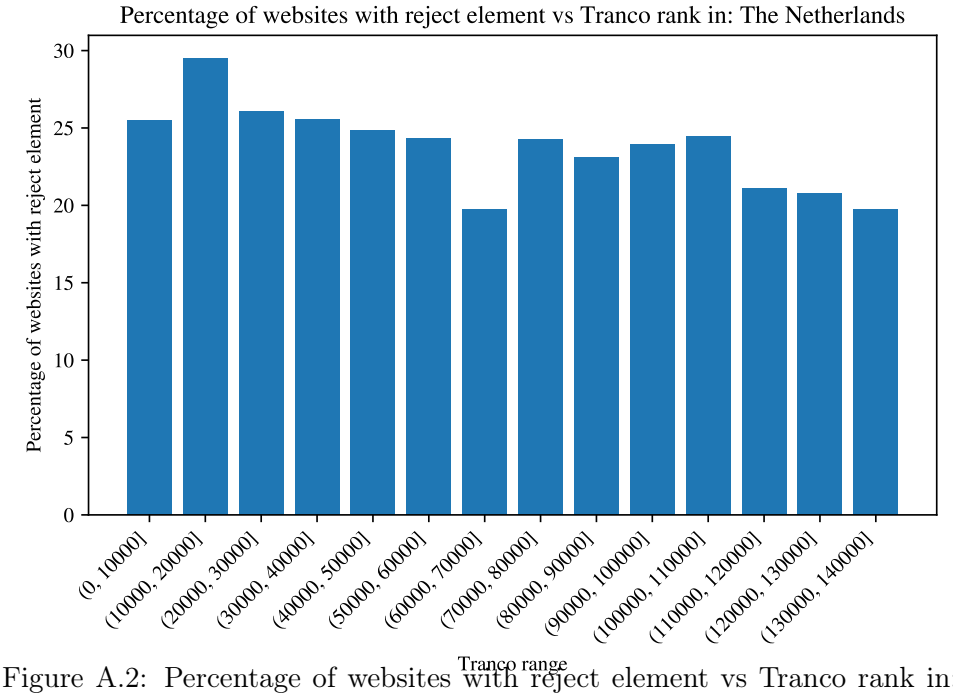

Figure A.2: Percentage of websites with reject element vs Tranco rank in: The Netherlands

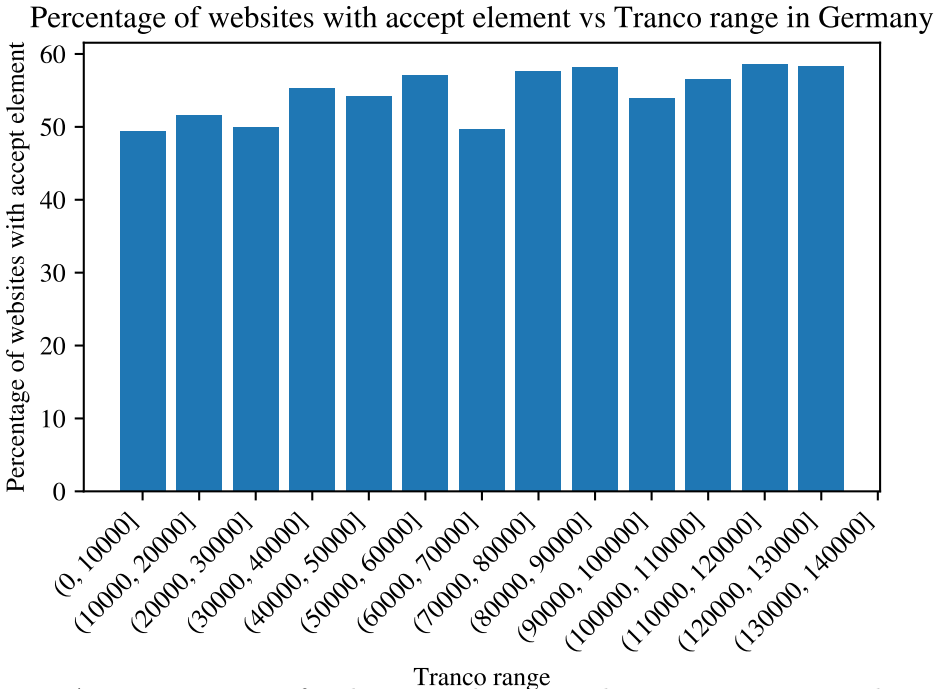

Figure A.3: Percentage of websites with accept element vs Tranco rank in: Germany

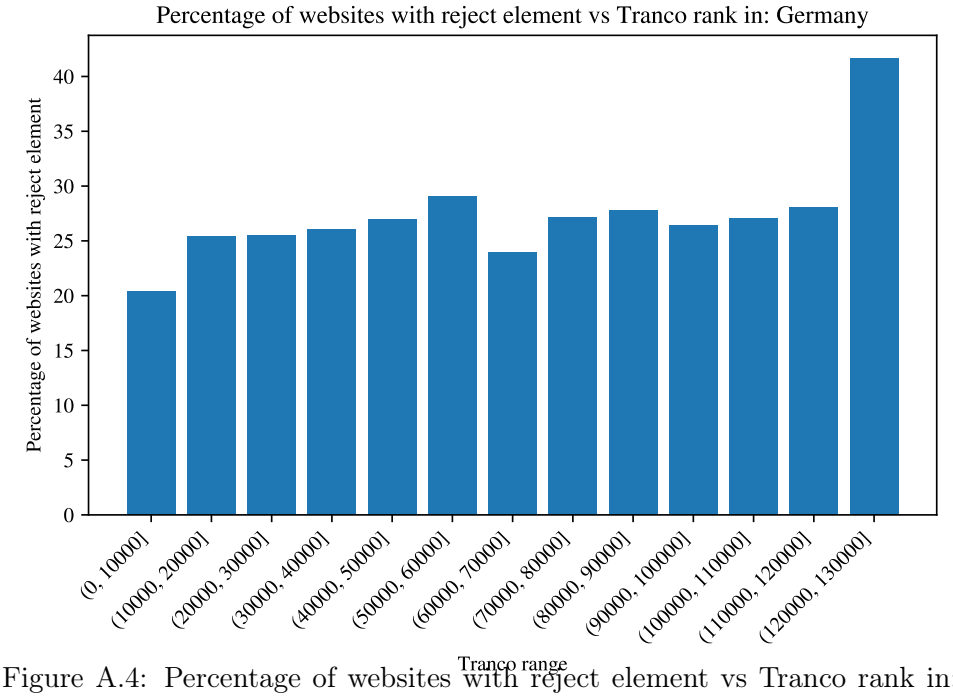

Figure A.4: Percentage of websites with reject element vs Tranco rank in: Germany

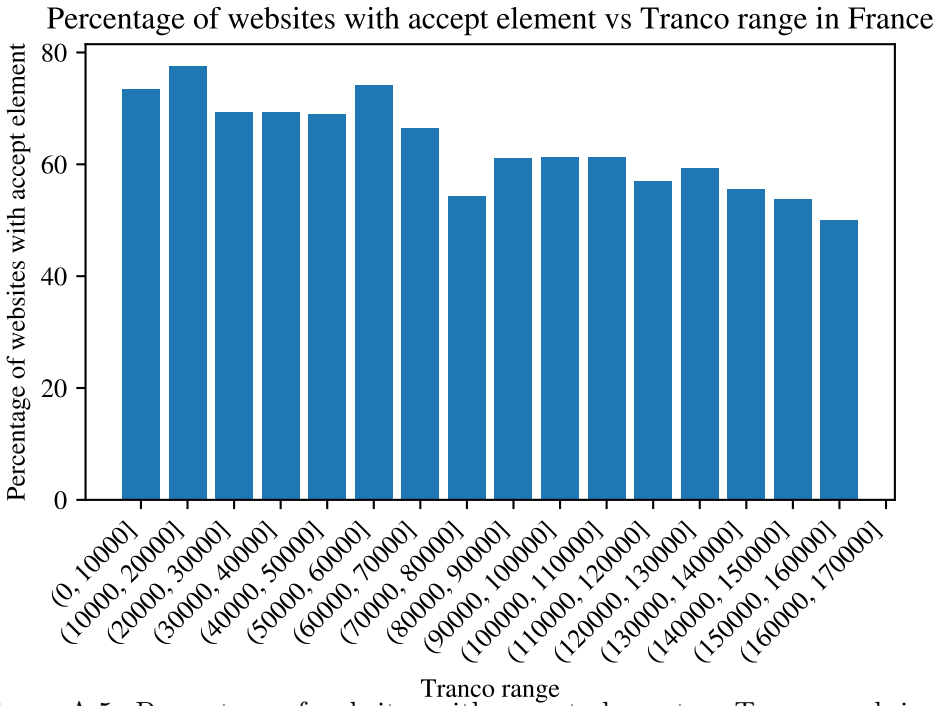

Figure A.5: Percentage of websites with accept element vs Tranco rank in: France

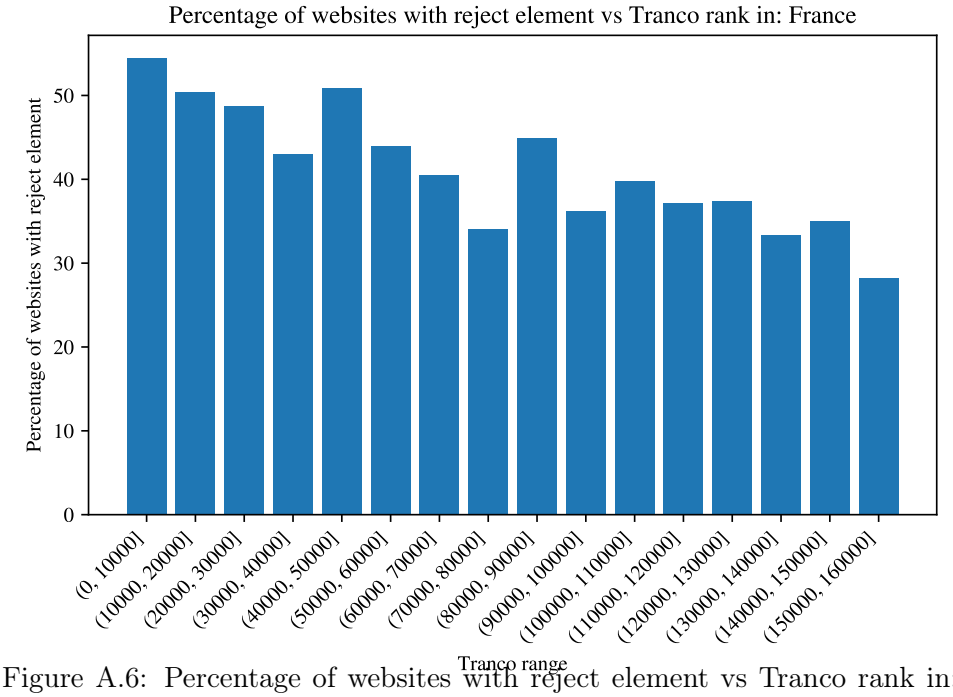

Figure A.6: Percentage of websites with reject element vs Tranco rank in: France

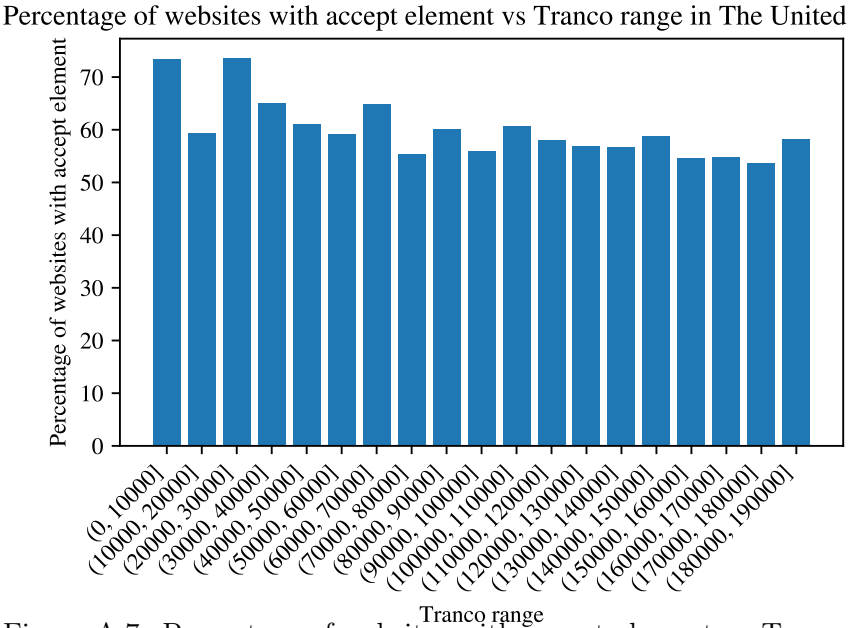

Figure A.7: Percentage of websites with accept element vs Tranco rank in: The United Kingdom

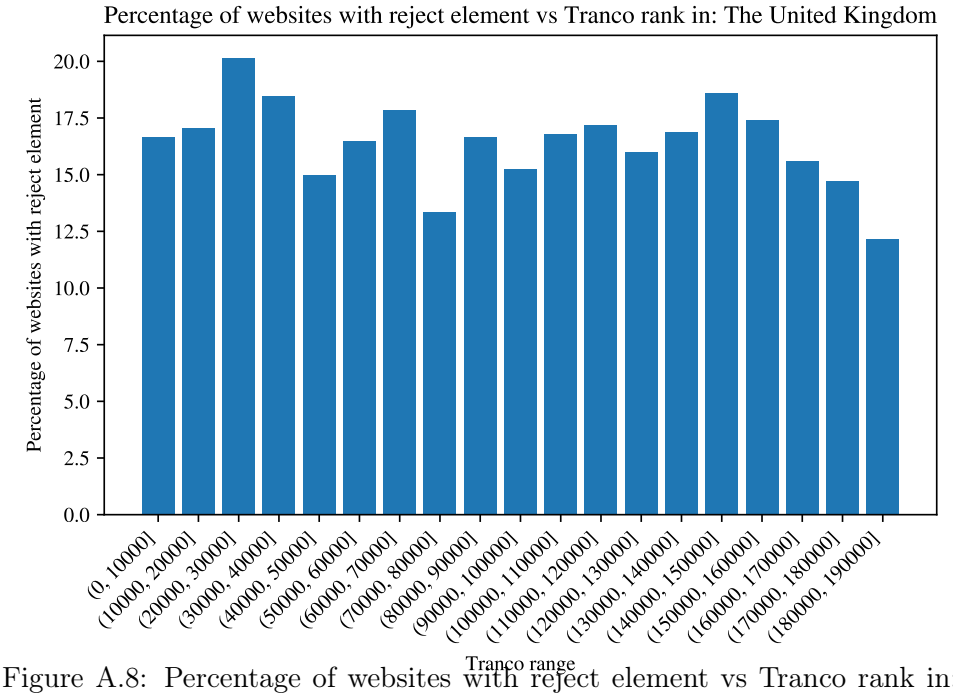

Figure A.8: Percentage of websites with reject element vs Tranco rank in: The United Kingdom

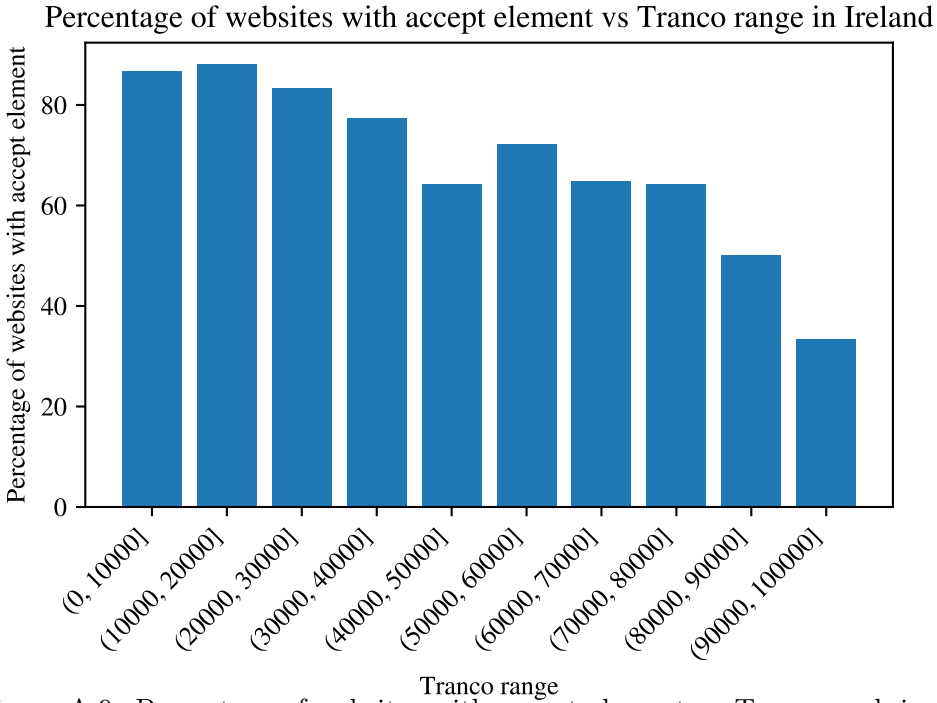

Figure A.9: Percentage of websites with accept element vs Tranco rank in: Ireland

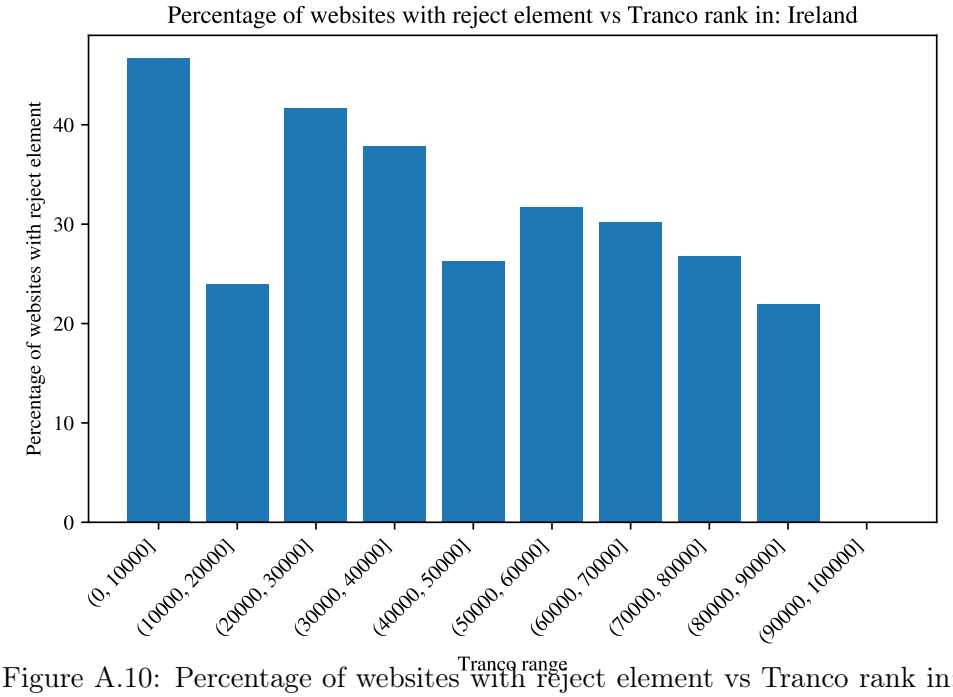

Figure A.10: Percentage of websites with reject element vs Tranco rank in: Ireland

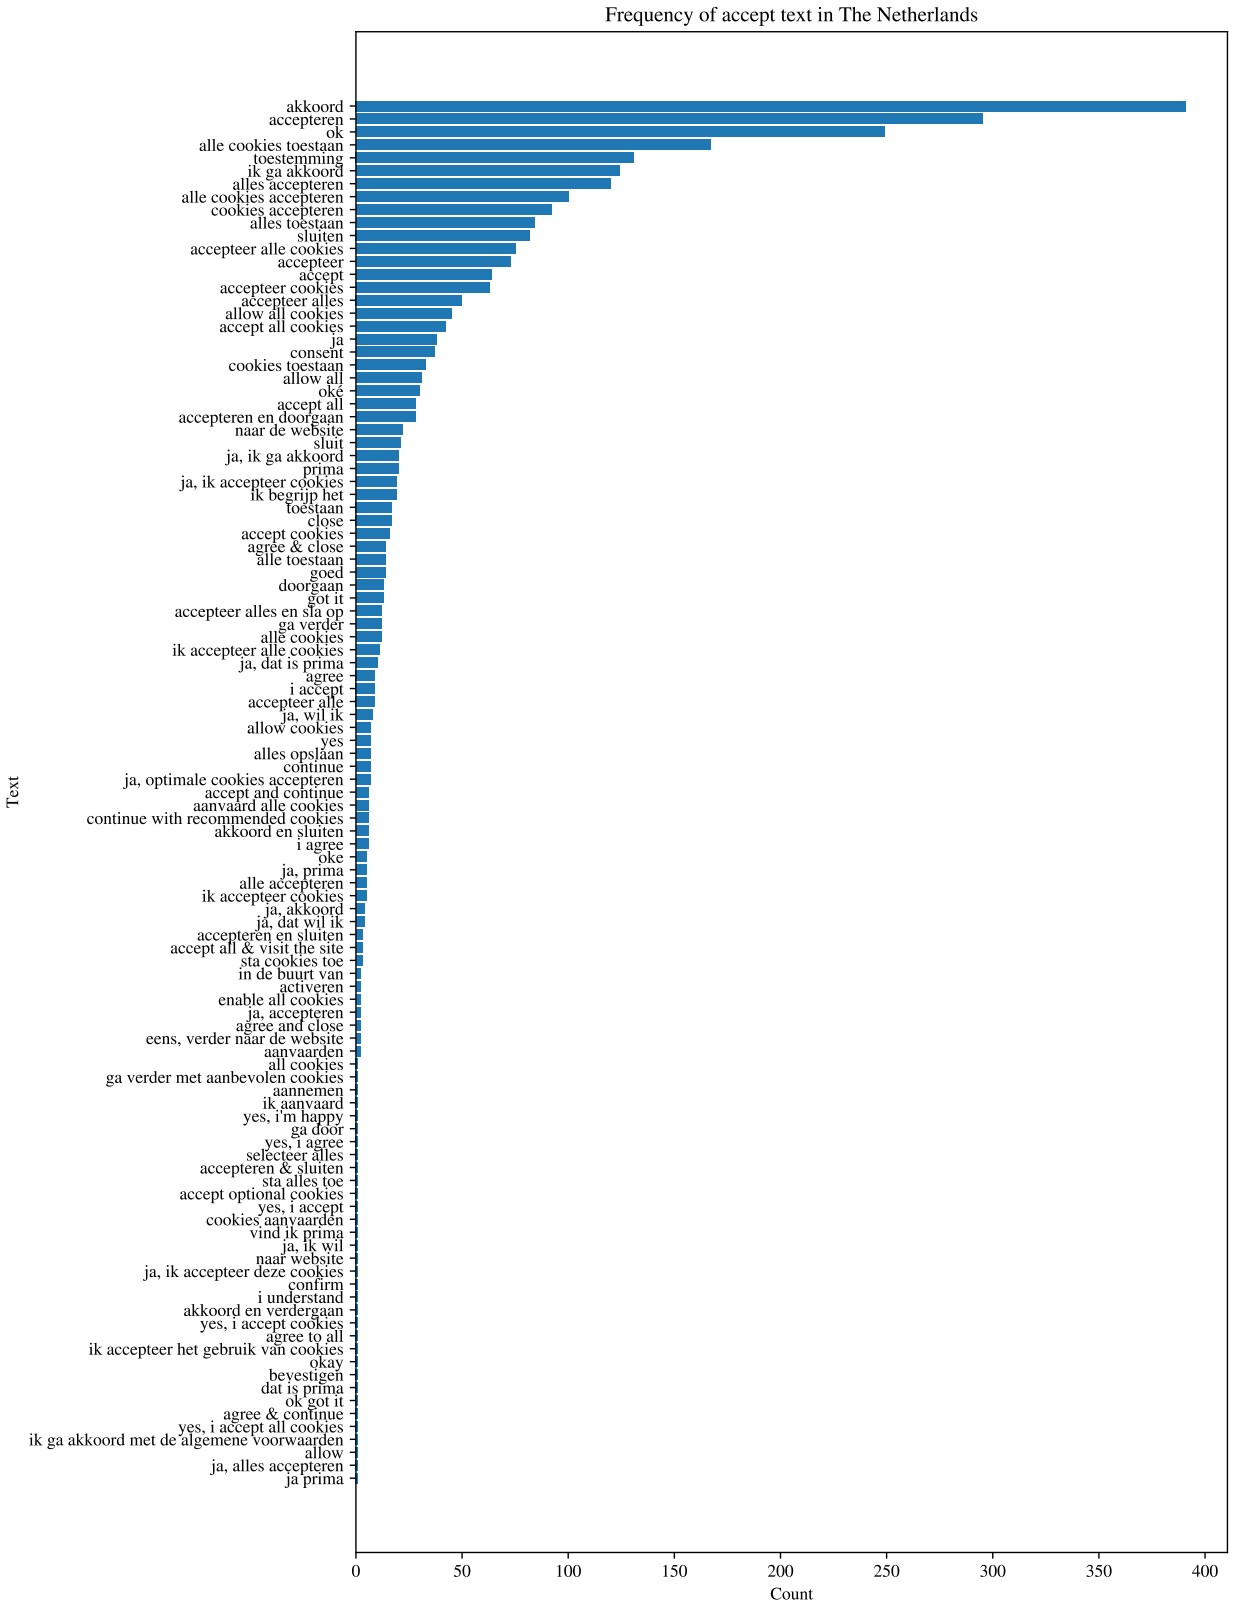

Figure A.11: Frequency of accept text in The Netherlands

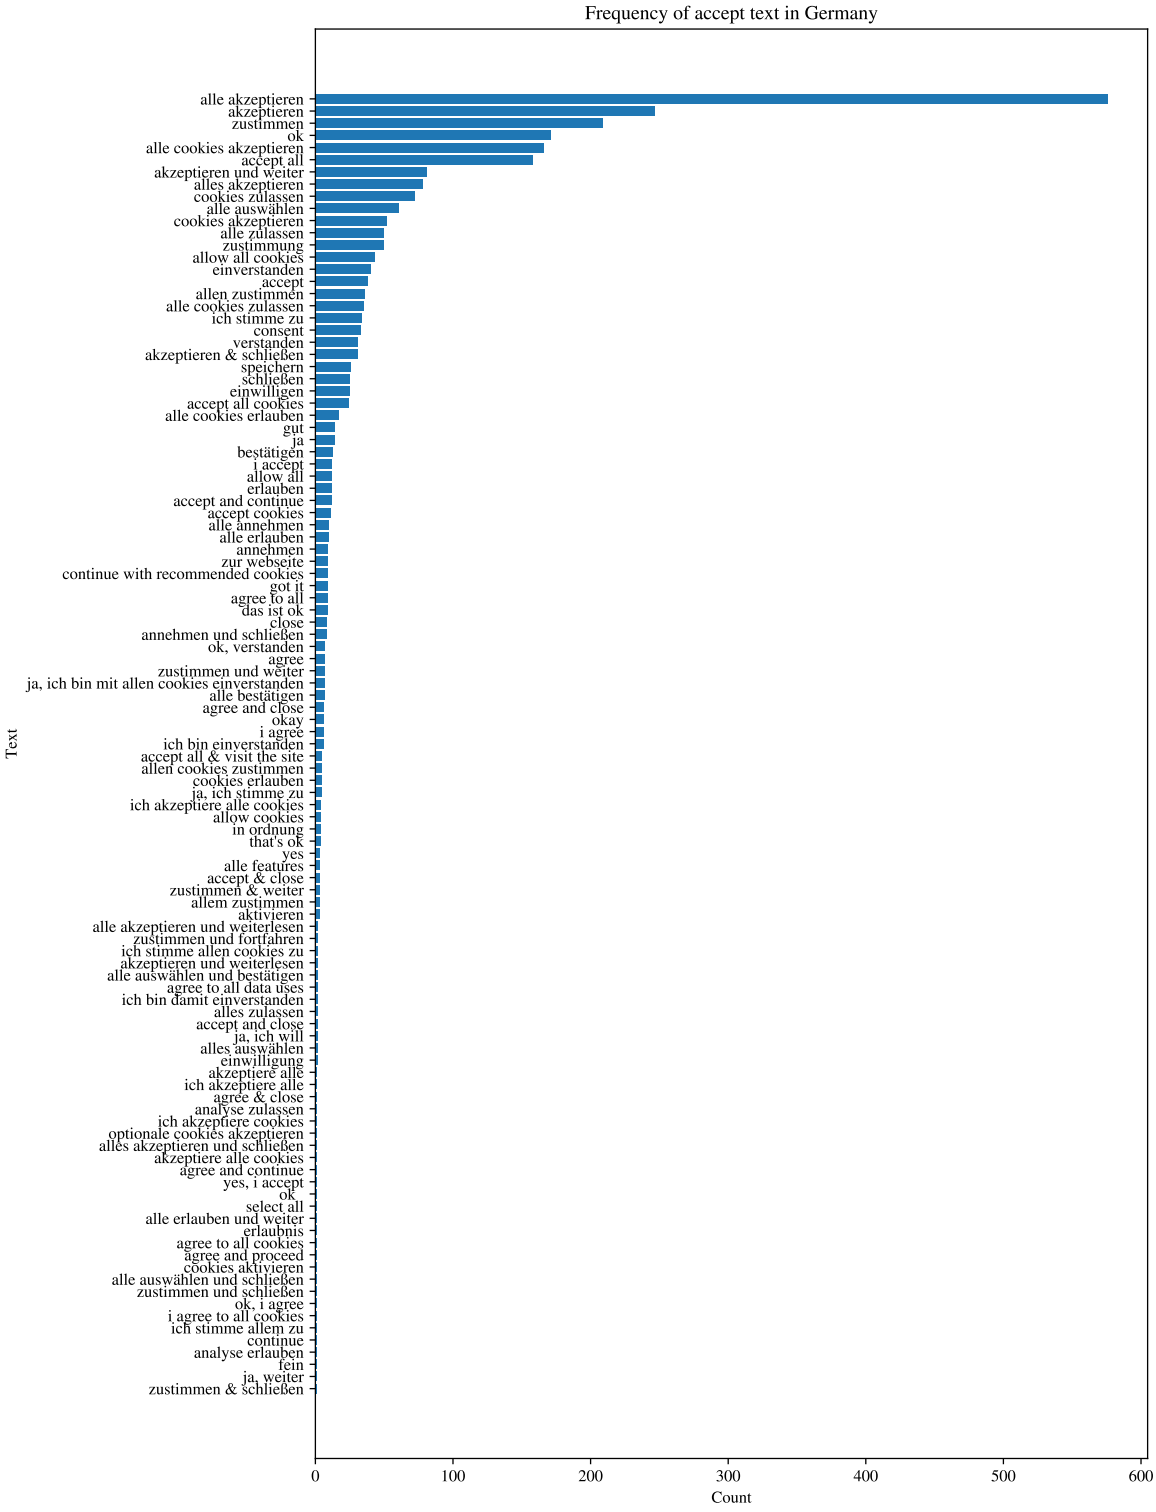

Figure A.12: Frequency of accept text in Germany

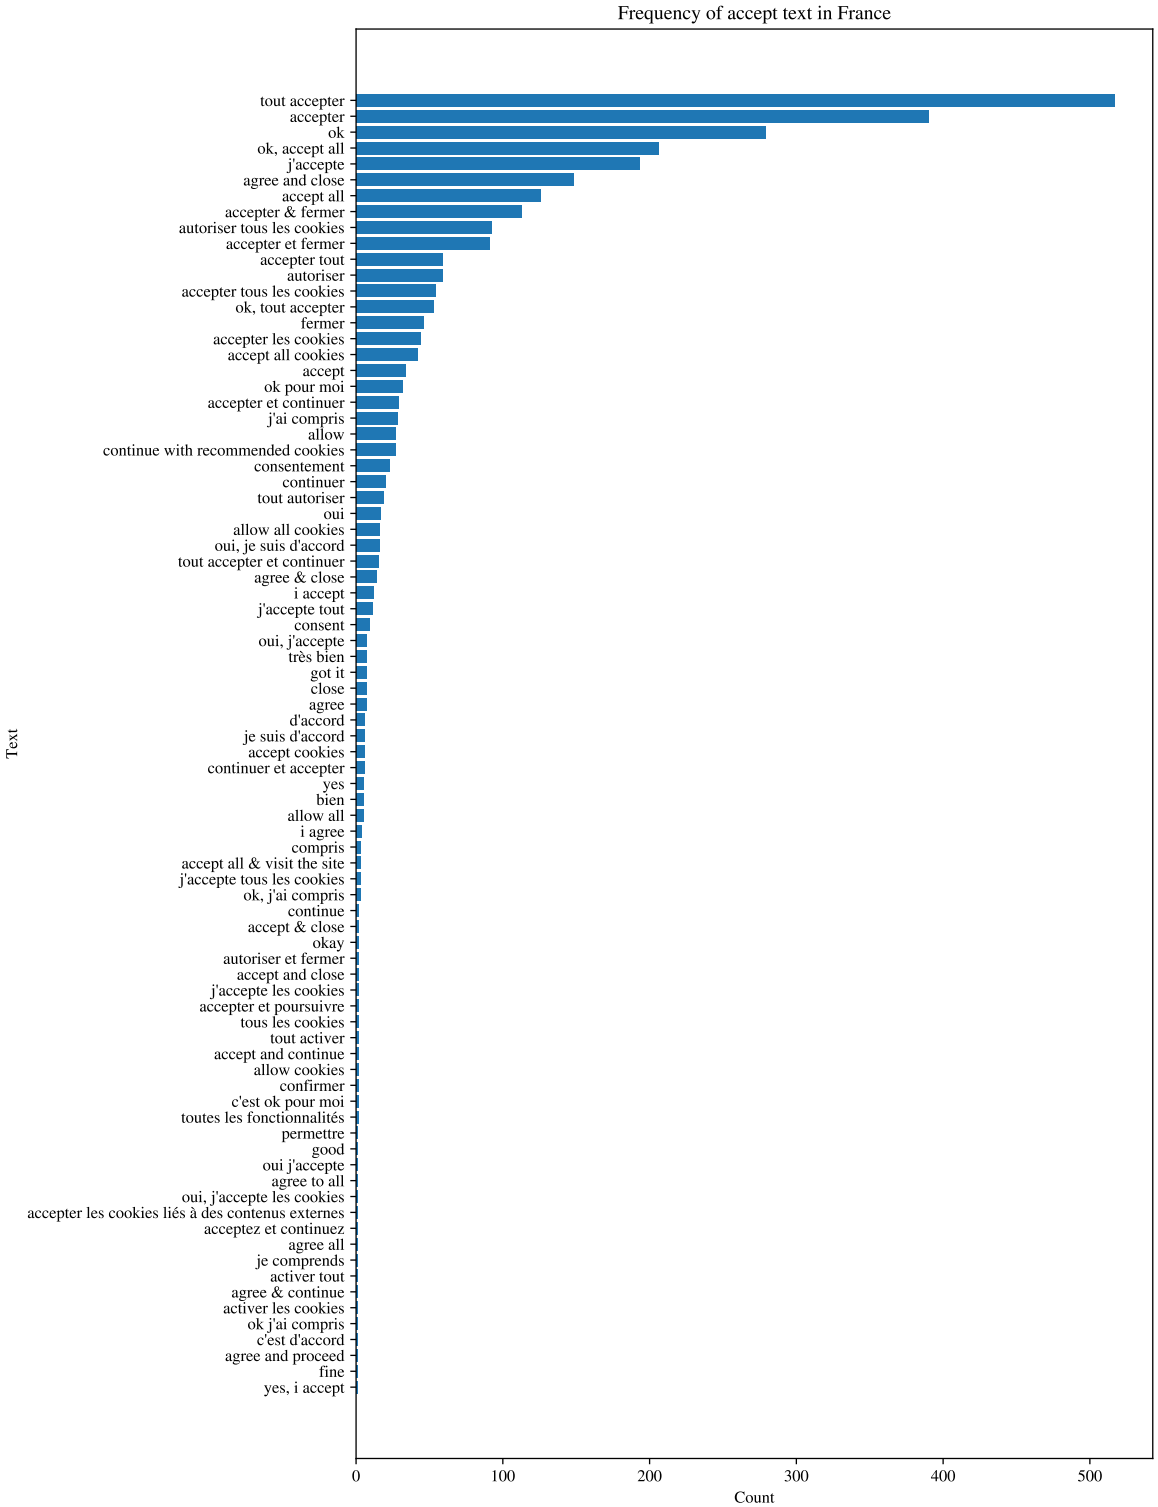

Figure A.13: Frequency of accept text in France

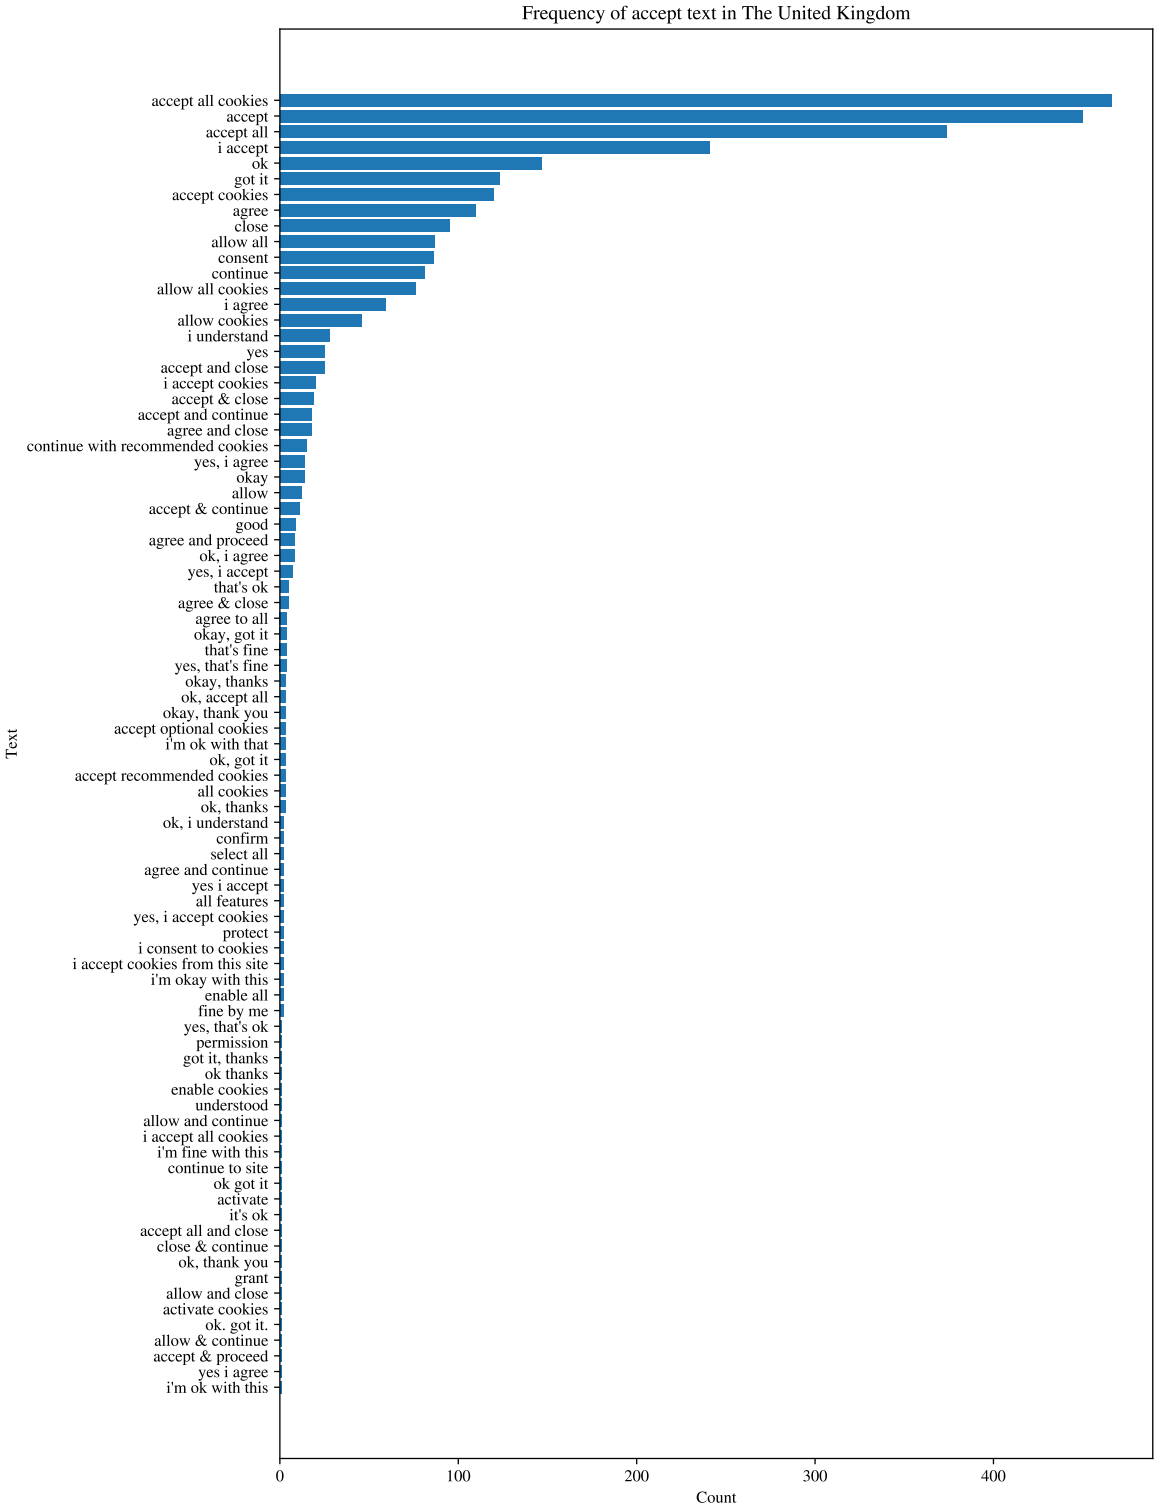

Figure A.14: Frequency of accept text in The United Kingdom

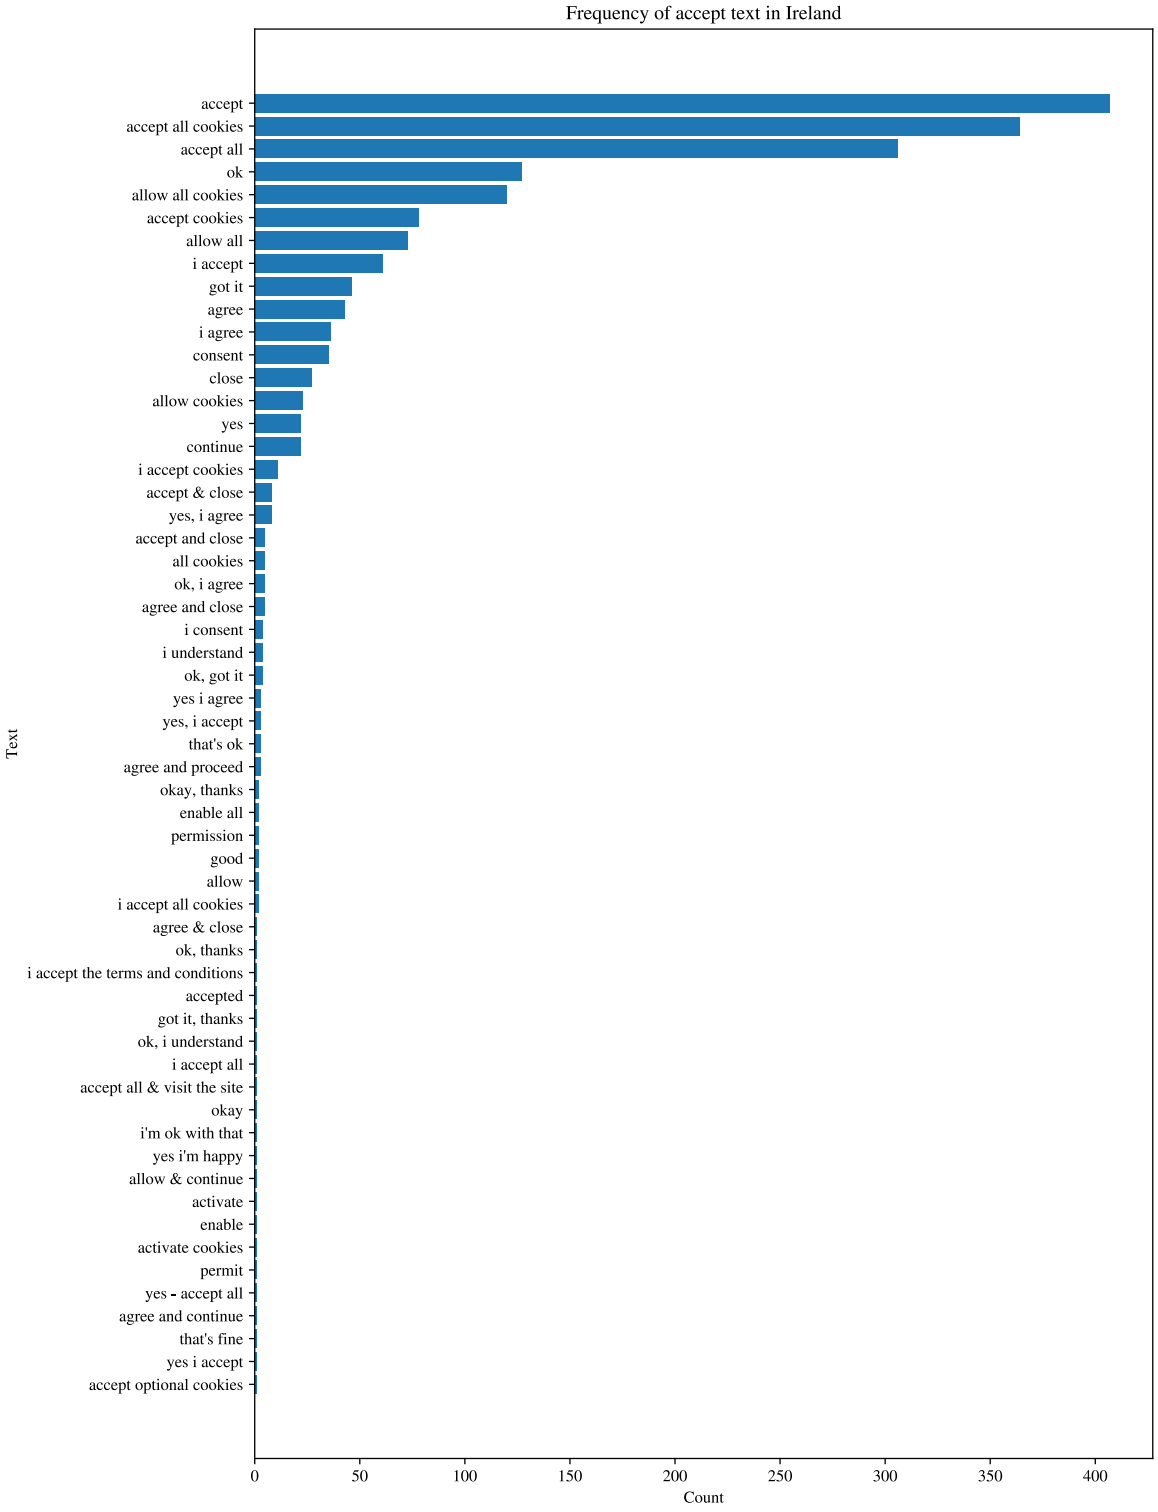

Figure A.15: Frequency of accept text in Ireland

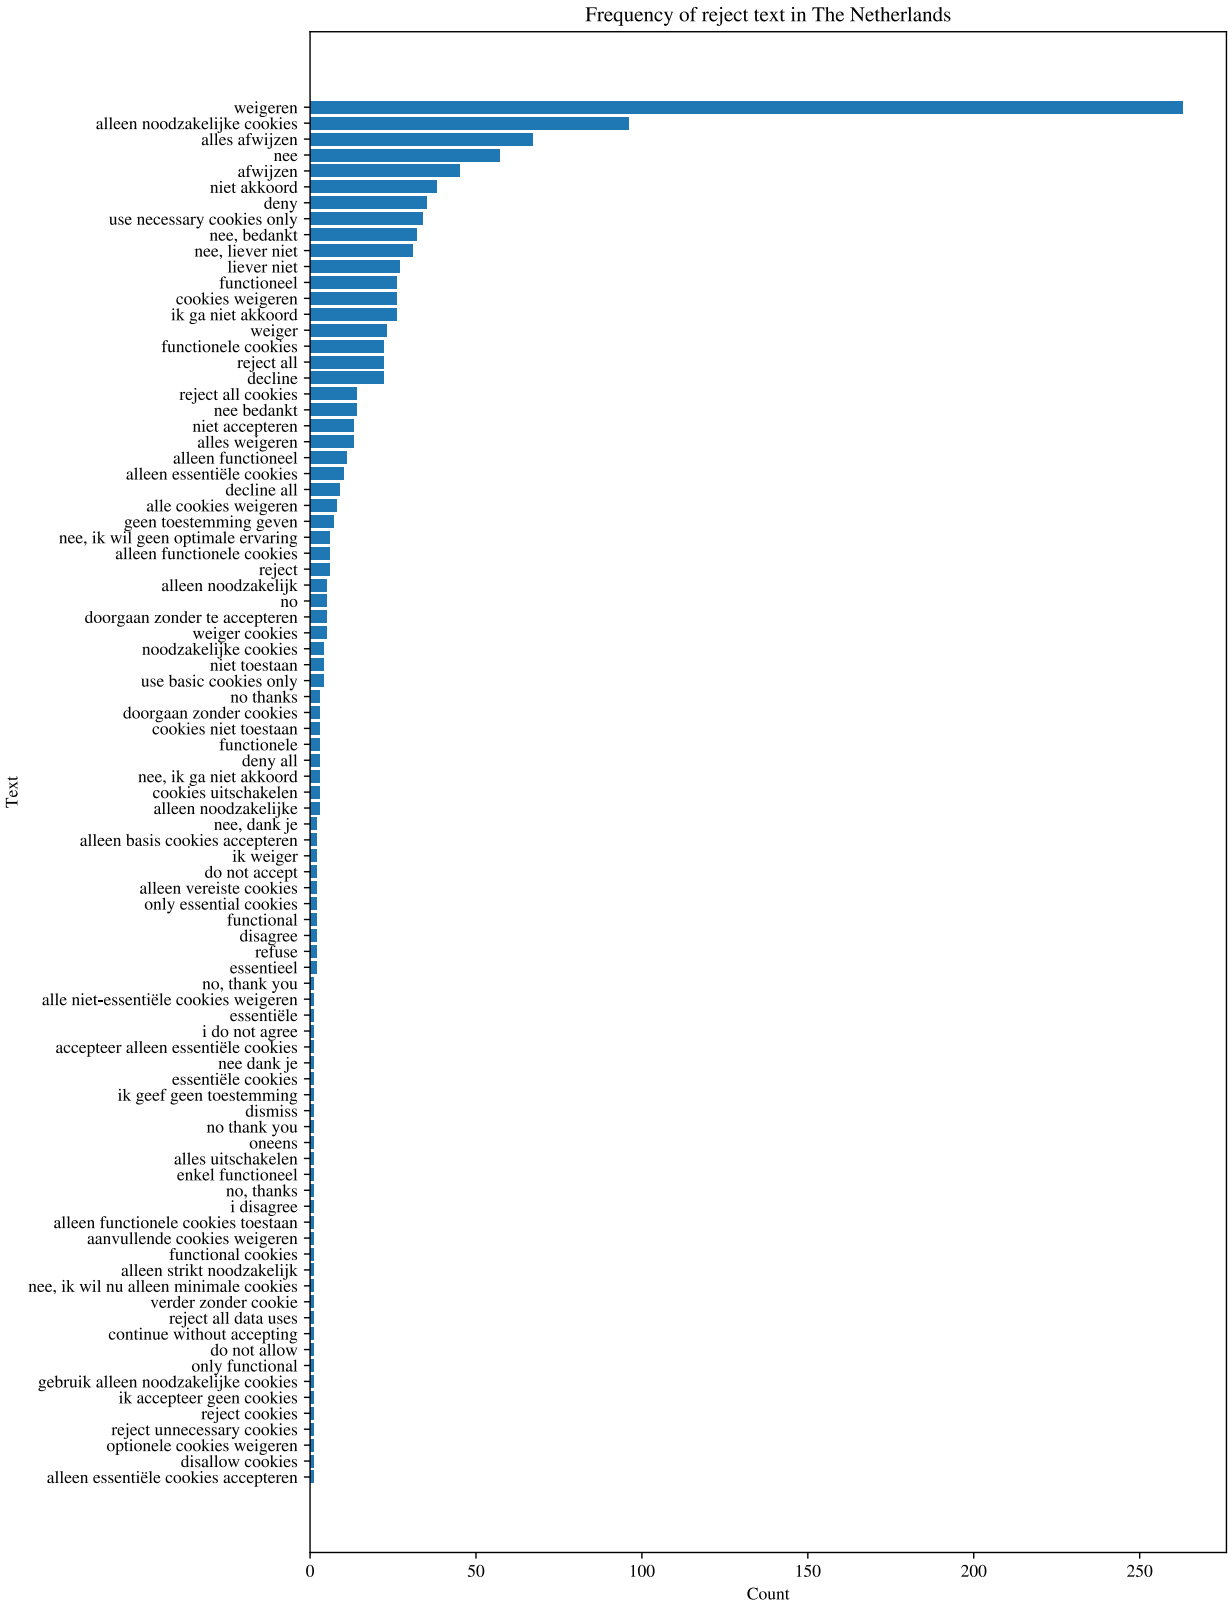

Figure A.16: Frequency of reject text in The Netherlands

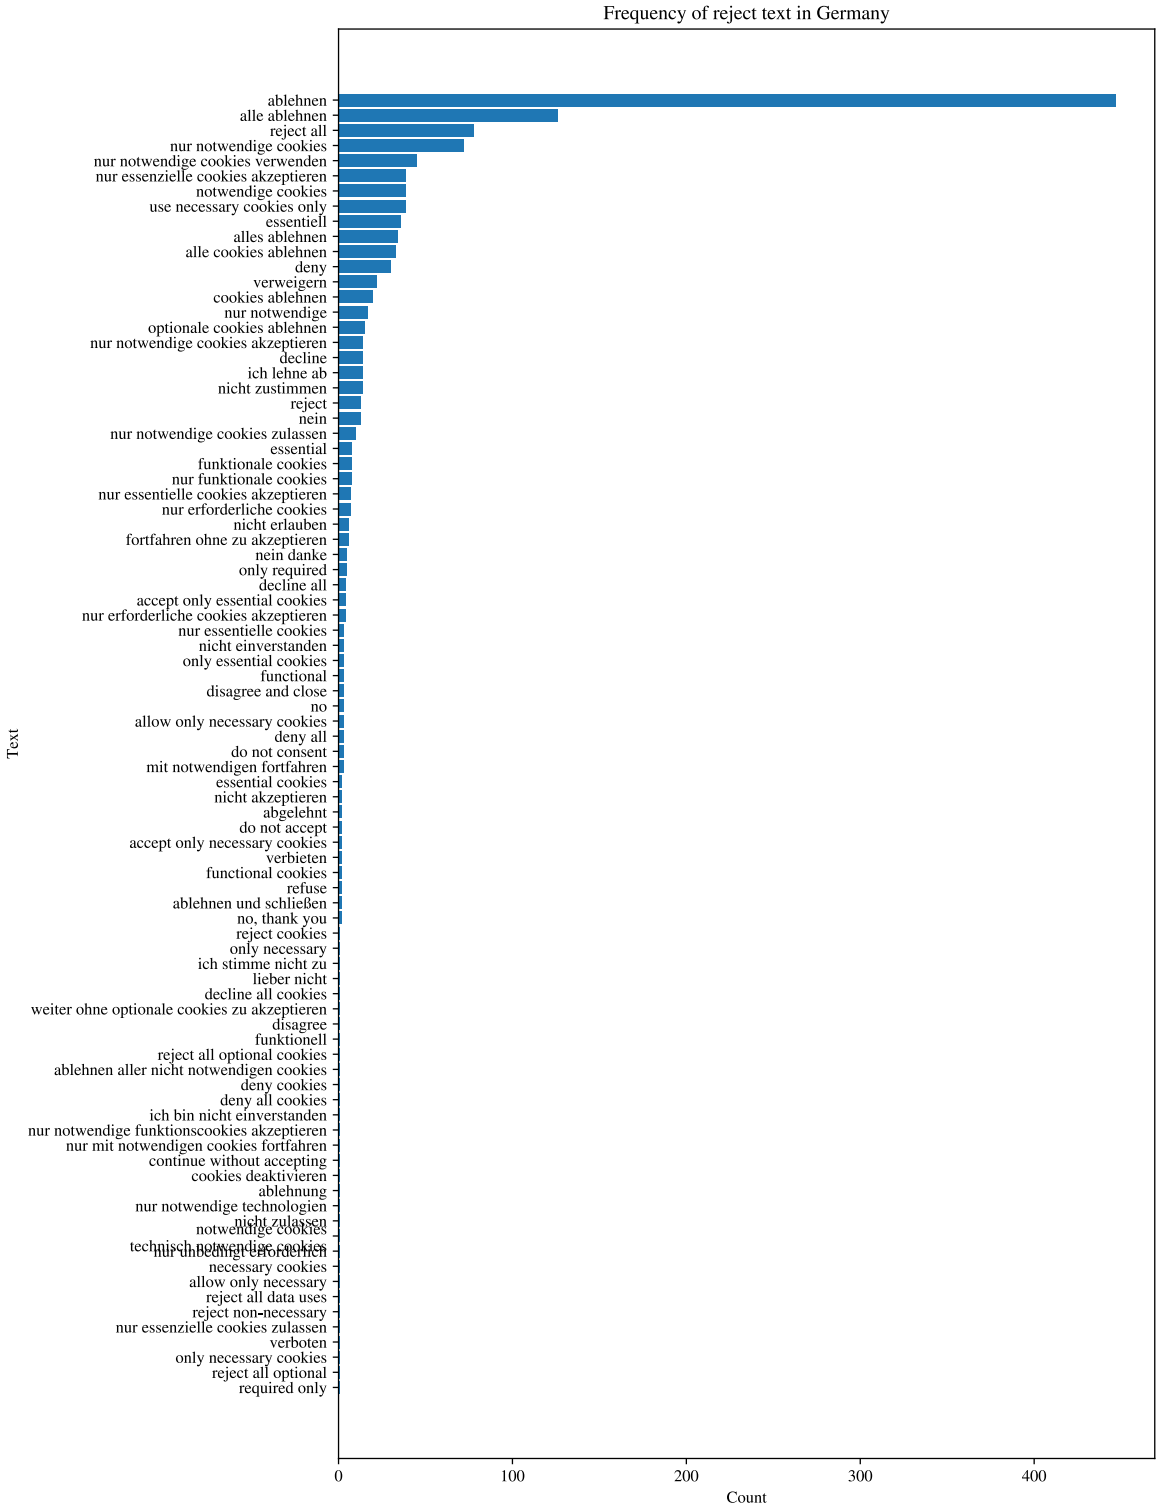

Figure A.17: Frequency of reject text in Germany

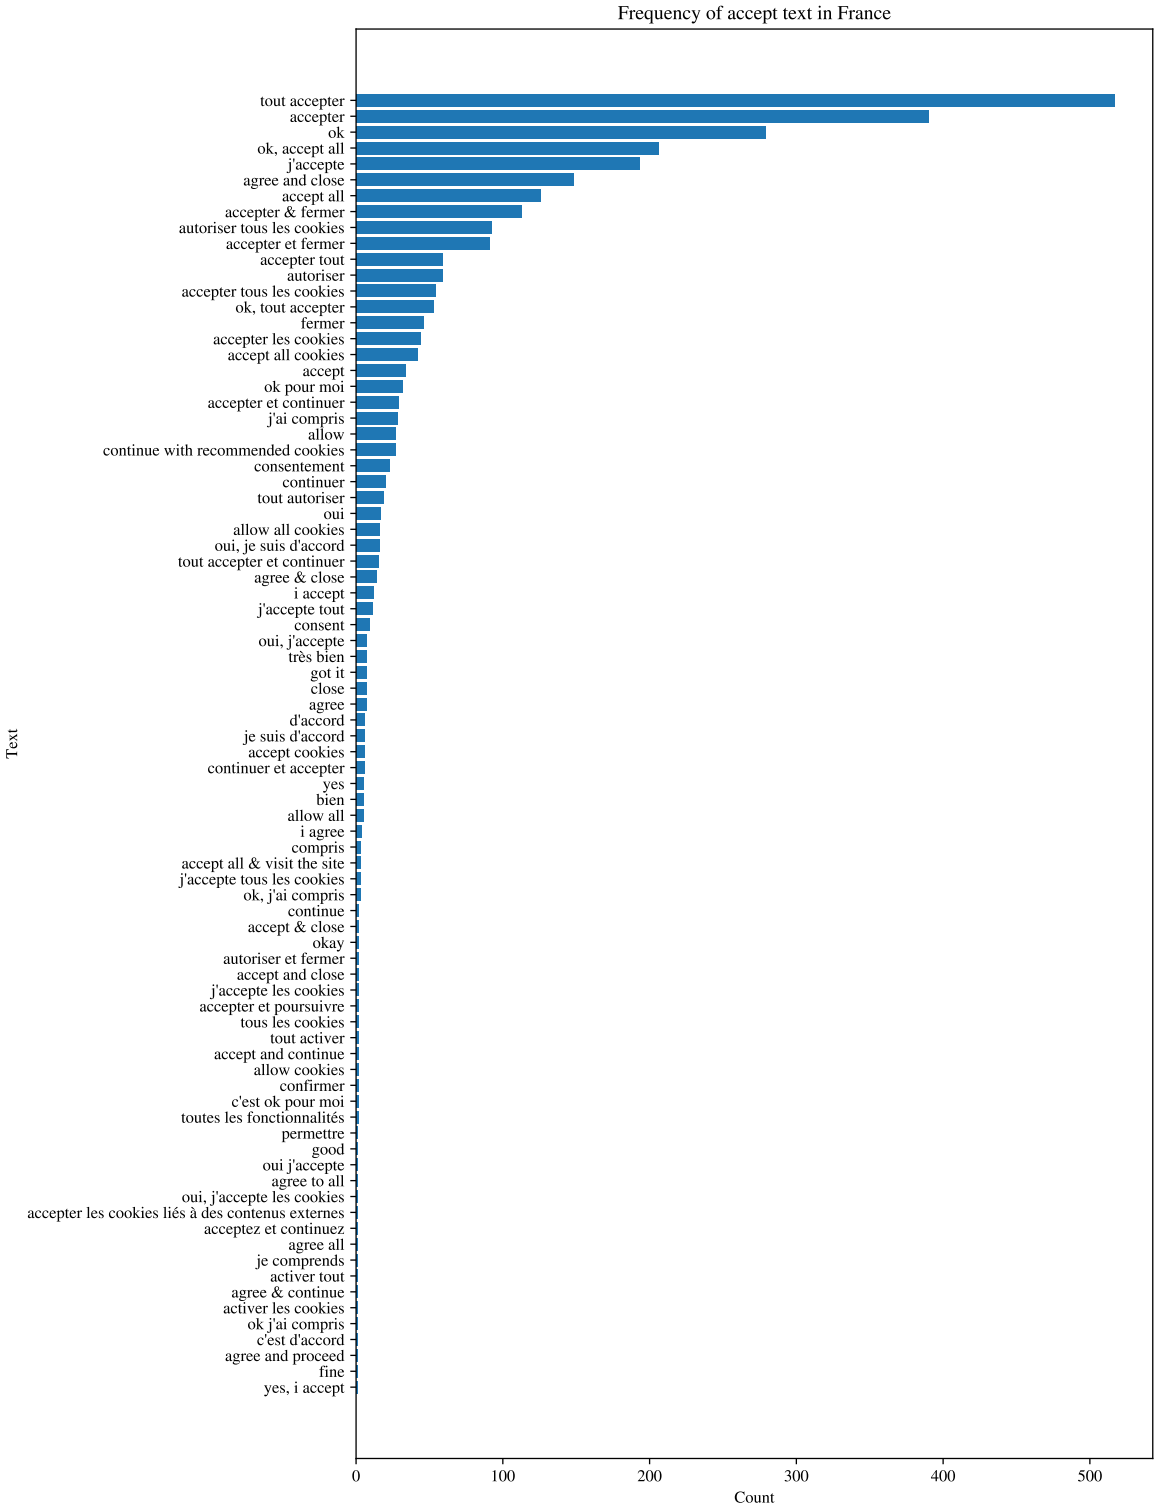

Figure A.18: Frequency of accept text in France

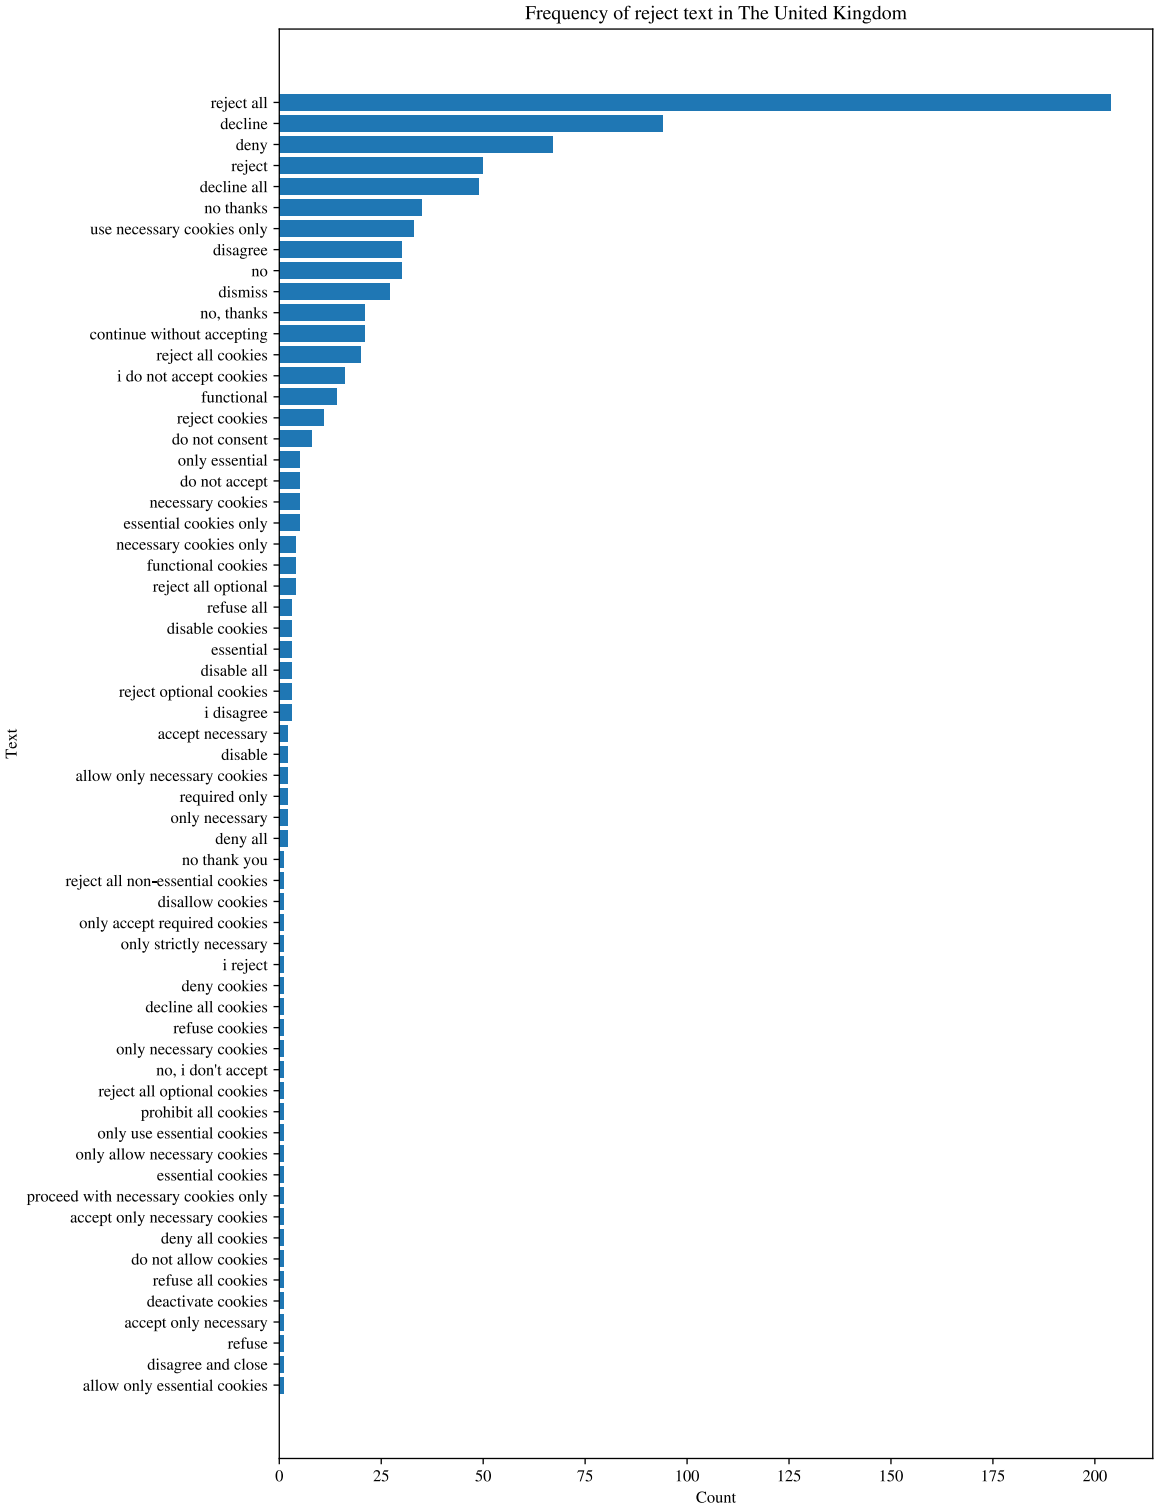

Figure A.19: Frequency of reject text in The United Kingdom

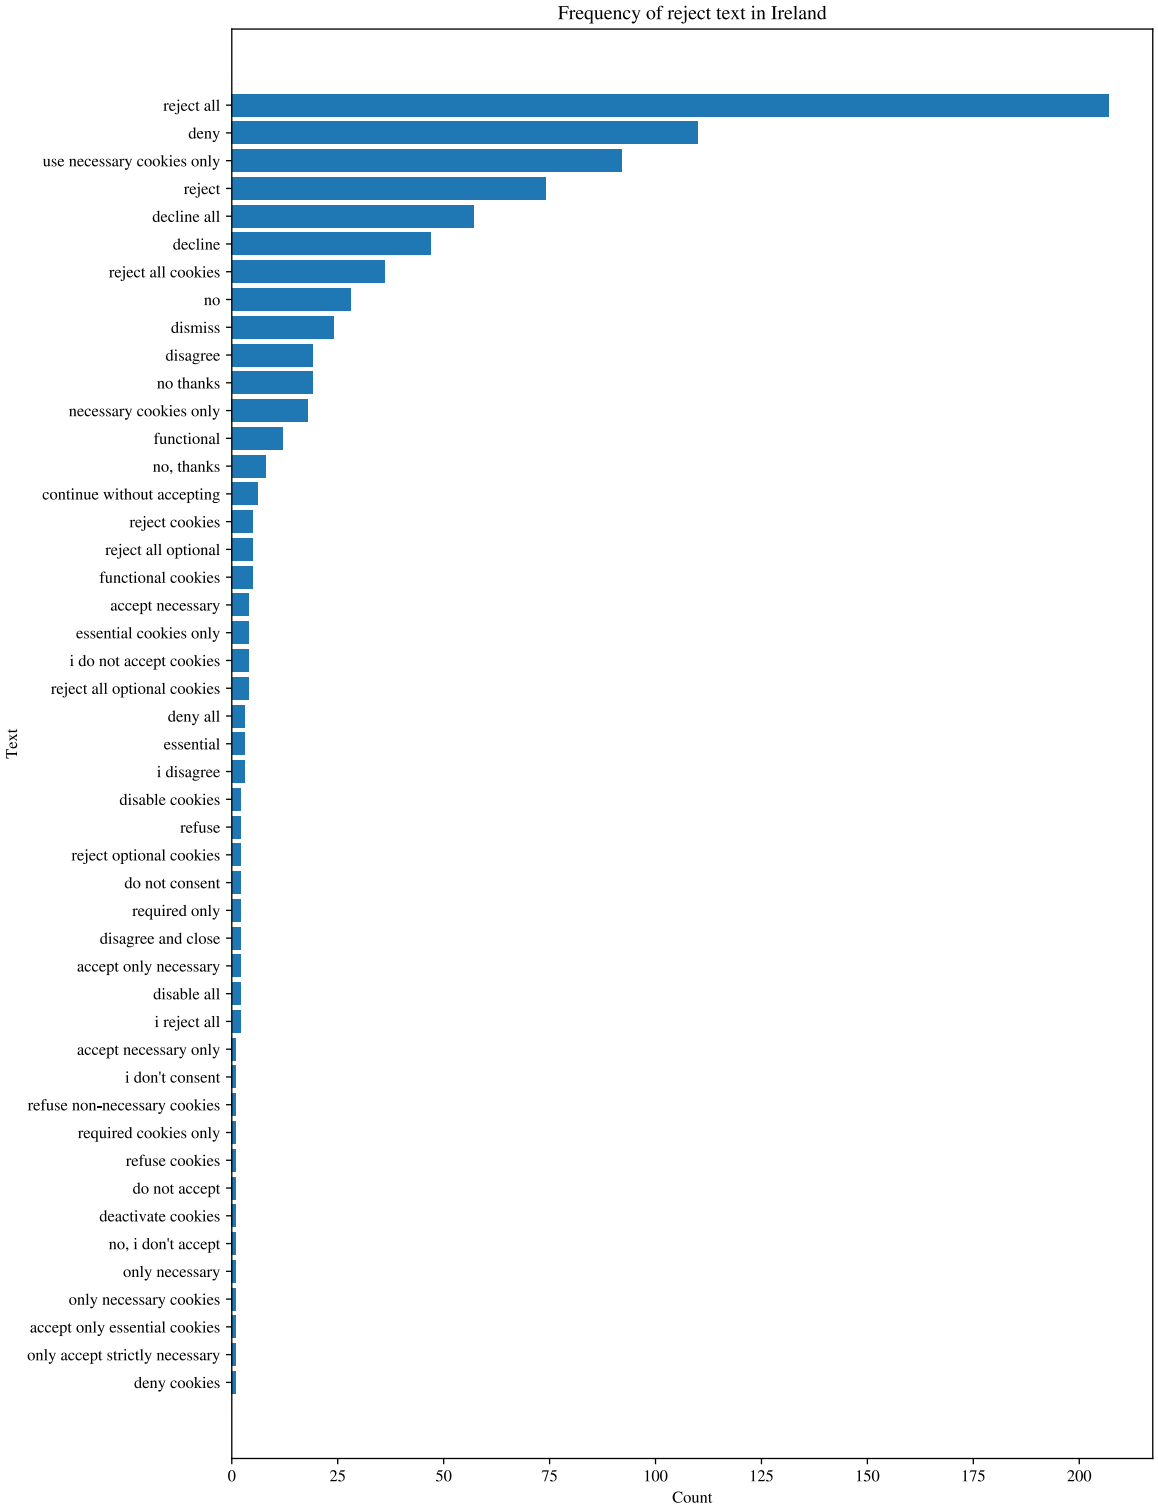

Figure A.20: Frequency of reject text in Ireland