

The Platform Governance Archive v1

A longitudinal dataset to study the governance of communication and interactions by platforms and the historical evolution of platform policies

Christian Katzenbach^{1, 2}, **Adrian Kopps**^{1, 2}, **João C. Magalhães**³, **Dennis Redeker**¹, **Tom Sühr**⁴, **Larissa Wunderlich**

¹ Centre for Media, Communication and Information Research (ZeMKI), University of Bremen

² Alexander von Humboldt Institut for Internet and Society, Berlin

³ Centre for Media and Journalism Studies, University of Groningen

⁴ Max Planck Institute for Intelligent Systems, Tübingen

Keywords: Platform Governance, Platform Policies, Terms of Service, Community Guidelines, Privacy Policies, Platform Law, Content Moderation, Digital Constitutionalism, Platform History, Archive

Dataset PGA v1

URL: <http://platformgovernancearchive.org/data/dataset-pga-v1-historical-dataset/>

DOI: <https://doi.org/10.17605/OSF.IO/8H2JG>

Licence: ODC-BY v1.0 <https://opendatacommons.org/licenses/by/1-0/>

How to Cite

Dataset PGA v1

Katzenbach, C., Kopps, A., Magalhaes, J. C., Redeker, D., Sühr, T. (2023). *Platform Governance Archive (PGA) v1*. [data set]. Centre for Media, Communication and Information Research (ZeMKI) University of Bremen. DOI: 10.17605/OSF.IO/8H2JG. URL: <http://platformgovernancearchive.org/data/dataset-pga-v1-historical-dataset/>.

Data Paper

Katzenbach, C., Kopps, A., Magalhaes, J. C., Redeker, D., Sühr, T., Wunderlich, L. (2023). *The Platform Governance Archive v1 – A longitudinal dataset to study the governance of communication and interactions by platforms and the historical evolution of platform policies*. Centre for Media, Communication and Information Research (ZeMKI), University of Bremen. <https://doi.org/10.26092/elib/2331>.

Cite a Single Policy Document

Name of platform. (Date of version). Name of policy. *Platform Governance Archive*. Direct URL.

Bremen / Berlin / Groningen, July 2023

Centre for Media, Communication and Information Research (ZeMKI)
University of Bremen

Extended abstract

Platform policies contain the spelled out rules about what is allowed and prohibited on a service. As such, they constitute both a normative framework as well as a means of public communication by platforms. Studying the evolution of the increasingly complex web of policies that platforms have developed can hence allow us to trace the emergence of a specific normative order, i.e. the ways in which platforms are governing user activities and public speech and communication dynamics, as well as identify how they have reacted to public controversies, political debates and legal regulation.

A major difficulty for studies on the historical evolution of platform policies, however, is the availability of past policies which is often needed for a thorough analysis, as the policies change quite frequently and even their names and locations often differ from the current version. Although platforms have become increasingly transparent about how and when they are changing their rules and have begun to offer public archives of the different historical versions of their policies, these archives often do not contain all of the past versions of a policy and relying on them entails trusting the platforms to provide complete information. Thus it remains hard to systematically study how the rules and norms of platforms have changed over time.

Our Platform Governance Archive (PGA) aims to address this need by providing a comprehensive and uniformly collected dataset of all of the historical versions of platform policies which does not rely on the platforms' own public records. While we are working on extending the scope of the archive to include more platforms and policies, the current dataset described in this paper contains all of the historical versions of three types of policy documents (Terms of Service, Community Guidelines, Privacy Policies) by four major platforms (Facebook, YouTube, Twitter and Instagram) in the time period from the inception of each policy until the end of 2021.

Our paper gives a comprehensive overview of the conceptual layout of the Platform Governance Archive and details the automated and manual processes of data collection and data cleaning, as well as our practical and theoretical challenges. Starting with how we define a relevant change to a platform policy, we lay out how we used the Internet Archive's Wayback Machine to identify past versions of platform policies, collect them, and then automatically and manually check for changes.

Specifically, we explain how we mapped the URLs of the selected policies and they have changed over time, putting together a puzzle of how they were renamed and relocated. We then detail the automated scraping process of these URLs from the Wayback Machine as well as the automated diff-checking which we employed. The last step of the data cleaning consisted in a manual revision of the automatically identified versions based on our definition of a relevant change, which was necessary because a significant amount of data noise remained. The paper furthermore describes how the platforms' ways of displaying their policies have changed over time by increasingly turning them into interactive pages and multi-page documents, as well as how we addressed the data collection challenges that arose from this.

The paper furthermore provides an overview of the resulting v1 corpus the Platform Governance Archive which is a dataset consisting of 354 policy documents with a total of 6,036 pages. By detailing the structure of our data repository on Github, we offer a guide on how to access and work with the data. We furthermore describe the characteristics and details of each platform and policy type to account for the fact that each of them have undergone a specific historical development.

Lastly, our paper also presents a structural analysis of some of the general trends and patterns which are visible in the dataset over a time period of up to almost two decades on the document level. Using a quantitative analysis, we analyse how the change frequency and the character count of each platform policy has developed over time. A comparative visualisation of these findings allows us to show how the extent of the policies has grown over time, to identify periods of high growth and frequent changes and to draw comparative conclusions about the four different platforms.

The Platform Governance Archive aims to be a resource for researchers, journalists, policy-makers, platform operators, activists, and other stakeholders as well as the general public. By offering both a comprehensive dataset and an accessible interface, we aim to offer and continue to develop this resource to enable research and public debate on the historical evolution of platform policies in order to trace down changes, to identify characteristic periods of isomorphic policies, to measure influencing factors, and to understand how specific debates, events, and legislation have influenced and manifested in platform policies.

Table of Content

1 Introduction	4
2 Context: Studying Platform Governance and Platform Policies	4
3 Introducing the Platform Governance Archive (PGA)	6
4 Methods of Data Collection and Data Cleaning	8
4.1 Building URL timelines	9
4.2 Automated scraping from the Wayback Machine	11
4.3 Automated checking for changes	12
4.4. Manual revision of the automatically identified versions	13
4.5. Collecting interactive pages and Multi-page documents	14
5 The PGA v1 Corpus	15
5.1 Project Website and Project Elements	15
5.2 Data Repositories: Folder Structure and Naming Conventions in the Corpus	16
5.2.1 /pga-corpus	16
5.2.1 /pga-workbench	17
5.3 Characteristics and Details of the Platforms and Policies	17
5.3.1 Facebook	17
5.3.2 YouTube	19
5.3.3 Twitter	20
5.3.4 Instagram	21
6 Trends in the Evolution of Platform Policies	22
6.1. Character count	22
6.1. Change frequency	25
7 Challenges and Limitations	28
8 Conclusion and Roadmap	29
Funding Note	30
Conflicts of Interest	30
References	30

1 Introduction

Over the past two decades, internet platforms have emerged as key social and technical institutions for communication and interactions. The ways in which they govern user activities has tremendous consequences for how our increasingly digital society is organised. Growing controversies around issues such as hate speech and misinformation point to the role and responsibility of platforms for the regulation of public speech and communication dynamics. Platforms have reacted by gradually abandoning their self-positioning as neutral technical intermediaries and by formulating increasingly extensive sets of rules that govern what kind of content and conduct are allowed or prohibited on their services.

While these platforms have offered more and more information about their moderation processes and decisions upon mounting public and policy pressure, it remains remarkably hard to systematically understand and study the changing rules and procedures of platforms' content moderation. How have Facebook's policies regarding misinformation changed over the time? When did Twitter include the line that it is the company's "purpose to serve the public conversation" in their "Twitter Rules"? And when did Facebook remove a provision that prohibits the posting of "false information" from their Terms of Service? Which cultural and political norms are embedded in the evermore extensive rulebooks that social media platforms have formulated? And which patterns and trends run through the historical evolution of these policies?

Starting from such questions we have built an extensive dataset of policy documents of major social media platforms that we call the *Platform Governance Archive (PGA)*. While we are working on successively extending the list of included platforms, the current dataset documented by this data paper consists of documents by four major platforms (Facebook, YouTube, Twitter and Instagram) and three types of policy documents (Terms of Service, Community Guidelines, Privacy Policies), with the exception of YouTube's Community Guidelines. Documents were collected and curated through a combination of automated and manual processes. This paper describes this hybrid data collection process, provides an overview of the dataset and of dynamics over a time period of up to seventeen years on the document level.

2 Context: Studying Platform Governance and Platform Policies

Platforms have become ubiquitous in today's online communication. In many sectors of society, platforms such as Ebay, Uber and Airbnb have emerged that serve to bring together sellers and buyers, or to solve other collective action problems. But the focus of this paper and the Platform Governance Archive (PGA) are social media platforms, which can be defined as "Internet-based, disentrained [asynchronous], and persistent channels of masspersonal communication facilitating perceptions of interactions among users, deriving value primarily from user-generated content" (Carr & Hayes, 2015, p. 49). Positioning themselves on a layer above other Internet-based applications) (cf. Gillespie, 2010, p. 351), social media platform companies have been able to scale up quickly, strongly based on their successful positioning in a sweet spot that allows them to "seek protection for facilitating user expression, yet also seek limited liability for what those users say" (Gillespie, 2010, p. 347).

The Responsibility Turn and the Growth of Platform Policies

Social media platforms for a long time, thus, presented themselves as governance intermediaries committed to free speech, with very limited rulebooks of what can be publicly communicated, and who would only be mandated by beneficial regulatory frameworks, such as Section 230 of the US Communications Decency Act (CDA) (Citron & Wittes, 2017) and the E-Commerce Directive 2001 in the European Union (Kuczerawy & Ausloos, 2015). These frameworks generally acquit platforms of direct liability – they are obliged to act only when being notified of potential violations and infringing content. In consequence, social media platforms have traditionally strongly upheld the stance of being neutral tech companies that merely facilitate free expression. This position had been famously expressed in Twitter's opening line of their Twitter Rules:

“Our goal is to provide a service that allows you to discover and receive content from sources that interest you as well as to share your content with others. We respect the ownership of the content that users share and each user is responsible for the content he or she provides. Because of these principles, we do not actively monitor user’s content and will not censor user content, except in limited circumstances described below.” (Twitter, 2009a, p. 1)

In consequence, social media platforms only got involved with a limited set of potential violations, most notably with regard to cases referring to privacy, and copyright infringements. Some platforms already published content policies in the late 2000s addressing undesired content such as violence, pornography, intimidation and abuse.¹ However, in retrospect, the scope of these policies was relatively limited. With the massive growth of their user base and time-spent on the platforms, social media started to be seen to affect political processes and outcomes (Barrett & Kreiss, 2019). They also proved to be a valuable source of almost unlimited personal information for intelligence agencies and a new form of “surveillance capitalism” (cf. Zuboff, 2019). As social media platforms grew over time, they achieved unprecedented notability in the political capitals of the world, including due to a number of scandals that were brought about by revelations of whistleblowers, journalists and activists. The resulting “techlash” in the US, Europe and beyond put additional pressure on these large platforms to create elaborate content policies and procedures to moderate illegal or undesired content. In this “turn to responsibility” (Katzenbach, 2021), platforms necessarily further abandoned their neutral position with regard to content that helped them to ascend in the first place. Substantiated by debates on misinformation, hate speech and violence in social media, there is now a “broad consensus that platforms have responsibility for the content and communication dynamics on their services” (Katzenbach, 2021, p. 3).

This turn has a noticeable effect on the self-positioning and the public ruleset of platforms. Twitter, for example, has changed its opening statement significantly. Since 7 June 2019, it reads:

“Twitter’s purpose is to serve the public conversation. Violence, harassment and other similar types of behaviour discourage people from expressing themselves, and ultimately diminish the value of global public conversation. Our rules are to ensure all people can participate in the public conversation freely and safely.” (Twitter, 2019, p. 1)

Generally, platforms have developed substantial, and more or less elaborate rulesets on content for their sites. These rulesets are increasingly documented on public pages and even separate websites.² Today, these written policies of large social media companies represent a comprehensive and powerful mode of governing communication and media on the internet. That online communication is governed in such a way by private actors rather than public entities may seem “lawless” in its current state due to a lack of legitimacy to rule (cf. Suzor, 2019). The facticity of what scholars have described as “platform law” (Bygrave, 2015; Kaye, 2019) remains: “digital intermediaries govern the Internet” (Suzor, 2019, p. 168). This is even the case where platform policies are not fully communicated, either because they are not yet codified (as in the early days of content moderation), or because the rules on censorship are intertwined with powerful sovereign interests (as, for instance, in the case of Chinese social media platforms), or because there are commercial interests that trump the rules that apply to everyone else leading to separated internal rulebooks (as in the recently revealed separate content moderation for celebrities’ content on Meta’s platforms)³. The lack of transparency of platform policies requires academics, policymakers, activists and journalists to remain vigilant and to continue pushing platforms to be as transparent as possible. Academic researchers can contribute substantially here to better understand how platform policies govern the Internet, how they change and how different platforms take different approaches to content governance.⁴

¹ Cf. eg. Facebook’s early “Content Code of Conduct” in 2007 (Facebook, 2007) and “The Twitter Rules” in 2009 (Twitter, 2009a).

² E.g. Meta’s Transparency Center (Meta, 2023b).

³ A 2021 leak, the so-called ‘Facebook Files’, entailed information about Meta’s XCheck program, which in late 2020 shielded at least 5.8 million important and celebrity users from the content moderation procedures applied to other users (Horwitz, 2021).

⁴ While content policies are the main focus of this portrait of platform governance, privacy policies and other material such as general terms of service are important go-to places to study how platforms govern, even if these are often more adaptive to increasingly strict data protection and privacy laws.

Studying Platform Policies

For these reasons, researchers have in recent years increasingly investigated platform policies as an instrumental source of information about how platforms govern communication, and how they relate and react to public and political debates. For example, scholars have highlighted how platforms respond to political processes by means of their policies, not only implementing regulatory changes but also reacting to the increasing pressure that they were facing with regard to their activities relevant to democratic processes such as elections (see e.g. Barrett & Kreiss, 2019; Katzenbach, 2021). Academics also engaged in comparisons between those Western platforms that have most strongly moved toward elaborate systems of content governance and those that promise to be “censorship-free” (Buckley & Schafer, 2022). Others have investigated specific platform policies, principles and sub-principles. Emillie de Keulenaar, João C. Magalhães and Bharath Ganesh have traced the evolution of Twitter policies on objectionable content (de Keulenaar et al., 2023). Rebecca Scharlach, Blake Hallinan and Limor Shifman have identified values and principles inscribed into platform policies (Scharlach et al., 2023). Eugenia Siapera and Paloma Viejo Otero have investigated hate speech policies and conclude that “Facebook’s content governance turns hate speech from a question of ethics, politics, and justice into a technical and logistical problem” (Siapera & Viejo Otero, 2021, p. 112).

What these studies have in common is that they rely on platform policies to analyse them based on a given theoretical framework. One major difficulty for such studies is the actual availability of past policies which is often needed for thorough analysis as they change quite frequently (Barrett & Kreiss, 2019; Katzenbach, 2021) – and even their names and locations often differ from the current version (cf. section 4.1). For instance, “Facebook as a platform undergoes a continual set of changes in its affordances, often without transparency and disclosure, as well as the seeming role of external pressure in driving them” (Barrett & Kreiss, 2019, p. 14). Archive data availability constitutes a key challenge for scholars when conducting research on platform policies. While platforms are increasingly transparent about how and when they are changing their rules and offer public archives of the different historical versions of their policies⁵, no legal requirement for archiving exists, putting researchers and others at the mercy of platforms’ voluntary reporting. It is time to create an independent archive for platform governance documents, driven by the academic community.

3 Introducing the Platform Governance Archive (PGA)

Against this background, we are developing the Platform Governance Archive (PGA) as a joint resource for researchers, journalists, policy-makers, platform operators, activists, and other stakeholders as well as the general public. It offers both a comprehensive dataset and an accessible interface to the content policies of major platforms over most of their lifetime, i.e. it includes all historical versions of platform policies back to the launch of each service (if available on the public internet). This allows for exploring the important and pressing issues of platform governance outlined above, most notably the identification of all the changes that occurred over time. Thus, we aim to offer and continue to develop this resource to enable research and public debate on the historical evolution of platform policies in order to trace down changes, to identify characteristic periods of isomorphic policies, to measure influencing factors, and to understand how specific debates, events, and legislation have influenced and manifested in platform policies.

Background and Related Initiatives

The PGA has emerged out of our data collection processes in different projects at the Alexander von Humboldt Institut for Internet and Society (HIIG), most notably the project “Platform Governance and Copyright”, funded by the European Union’s Horizon 2020 research and innovation programme as part of the ReCreating Europe consortium (Grant Agreement Number 870626).⁶ The project is

⁵ E.g. Meta’s Transparency Center (Meta, 2023b), Google’s archives of Privacy Policies (Google, 2023a) and Terms of Service (Google, 2023b) and Twitter’s Terms of Service archive (Twitter, 2023a)

⁶ Cf. the project page at

inspired by Kirsten Gollatz' pioneering work on the longitudinal collection and analysis of platform policies in her dissertation "Die private Governance der Plattformen zur Regelung grenzüberschreitender Kommunikation" at the University of Zurich (Gollatz, 2019). The PGA is now hosted by the Lab Platform Governance, Media, and Technology at the Centre for Media, Communication and Information Research (ZeMKI), University of Bremen. We collaborate with the Alexander von Humboldt Institute for Internet and Society (HIIG), the University of Groningen, and Open Terms Archive for the curation and collection of the policies.

With its offering, the PGA is part of a set of initiatives that give visibility and transparency to platform policies: the Princeton-Leuven Longitudinal Corpus of Privacy Policies (Amos et al., 2020) is a massive dataset of privacy policies of major website; the Open Terms Archive (2023) is tracking almost in real-time changes of policies across more than hundred major websites; the Bipartisan Policy Center (2022) and the Integrity Institute are tracking platform announcements on election-related policy changes in their Technology Platforms Election Database. The Shorenstein Center at Harvard Kennedy School's FB Archive is storing the "Facebook Files" originating from whistleblower Frances Haugen (Harvard Kennedy School, 2023). The Social Media Archive (SOMAR) at the University of Michigan is a dedicated repository for social media research data, mostly on content-level (The Regents of the University of Michigan, 2023).

Profile and Scope of the PGA

The specific profile of the PGA in this line-up of important initiatives is that we focus on the most important social media platforms that substantially reorganise social communication which gives their content policies massive power for the organisation of public and private communication, as well as democratic processes; (b) for this goal, we include platform content moderation policies, often called community guidelines, that are the key locus for contested policies such as drawing lines on hate speech and misinformation, in contrast to the legalistic Terms of Service archived by most other initiatives; (c) we balance scale and substantial quality checks by a combination of automated and manual collection processes which allows us to include the much more distributed content policies of community guidelines as well well to identify when exactly actual changes to the policies occurred (instead of simply trusting the changes and timestamps shown in the platform companies' own archives).

So what is the scope of the PGA? We aim to archive *principal policy documents* of *popular social media platforms*. Different academic disciplines and policy makers disagree on what social media platforms are (and which of them, due to their importance, take on a "gatekeeper" or "custodian" role). It is often easier to enumerate examples of social media platforms than to arrive at a conclusive definition – "although we know what social media are, we are not necessarily able to articulate why they are what they are" (Carr & Hayes, 2015, p. 46). Carr and Hayes definition of social media can provide a starting point for a logic according to which the PGA selects social media platforms for archiving: according to them social media platforms are "Internet-based, disentrained [asynchronous], and persistent channels of masspersonal communication facilitating perceptions of interactions among users, deriving value primarily from user-generated content" (2015, p. 49). Within that space, we prioritise platforms that are created rather for general purpose communication than for specific communication purposes such as dating or pornography. We also focus on platforms that give rise to substantial "networked publics" (boyd, 2010) rather than pure messenger services. Yet such lines are increasingly becoming blurry with services such as Telegram offering channels with massive numbers of recipients.

In this first iteration, the Platform Governance Archive includes four platforms: Facebook, Instagram, Twitter, and YouTube. But we are already working on including many more platforms that match the above characterised criteria. In principle, we aim to archive the content governance documents of the most-used social media platforms. These larger social media platforms arguably function as 'gatekeepers', e.g., the EU defines the largest platforms in its Digital Services Act (DSA) package (at least 45m active EU end users/month and at least 10,000 annual active EU businesses, see DSA Chapter 2, Art. 3, 2b), or as 'custodians of the Internet' in the words of Tarleton Gillespie (2018). In order to assess future inclusions, we work with the list and criteria compiled by Chand

<https://www.uni-bremen.de/en/zemki/research/current-research-projects/platform-governance-and-copyright-h2020> [4 June 2023]

Rajendra-Nicolucci and Ethan Zuckerman (2021) of the 100 largest social media companies. To arrive at a demarcation of social media, they ask whether a platform is “a site with social features’ or is it ‘social first’?” and opt not to include websites that offer social features such as the abilities to post comments under articles or reviews for products (Rajendra-Nicolucci & Zuckerman, 2021). Based on data from Alexa, SimilarWeb and Wikipedia, they created a list of the most popular “social first” platforms. They exclude dating apps because they see them “more akin to platforms like Uber which operate two-sided marketplaces” and services like Zoom or Skype because they consider them “more of a utility, like the telephone”.

Platform	Documents		
Facebook	Terms of Service	Privacy Policy	Community Guidelines
Instagram	Terms of Service	Privacy Policy	Community Guidelines
Twitter	Terms of Service	Privacy Policy	Community Guidelines
YouTube	Terms of Service	Privacy Policy	–

Table 1: Set of platforms and documents included in this version of PGA

In terms of *policy documents*, we archive those documents published by platforms that describe their policies with regard to content on their services. This means we (aim to) include legal documents such as Privacy Policies and Terms of Service, as well as more prose-form documents such as Community Guidelines and content moderation principles. This genre of documents, explaining and legitimising content moderation and policy decisions, is undergoing dynamic development. During the time period under study in the PGA, especially Community Guidelines have evolved from short single-page-documents to complex multi-layered websites offering ever more information, details, and examples of how platforms take content-related decisions. This is also the reason why we were not able to include the Community Guidelines of YouTube in this version yet, because it has branched out in so many sub-pages and drop-down-menus that the collection was not feasible for our hybrid automatic-manual process (cf. chapter 4). For that reason, we anticipate that we will adapt and continue to develop this dimension of the PGA in the future. In contrast, we do not include API policies or other more procedural or technical legal documents as they do not cover substantially the kind of content policy decisions that are fundamental for social communications and that is the focus of the PGA.

In short, the PGA aims to document the services and platforms that are most influential with regard to social communication by archiving the public documents that explain how they govern and curate the communication and activities under their custody. We expect that this continues to be a key site of digital policy debate for the years to come as platforms and their governance are in the process of institutionalisation, and the platforms covered by the PGA will be those where content moderation issues are most salient and visible.

4 Methods of Data Collection and Data Cleaning

Data collection for the Platform Governance Archive (PGA) consisted of a multi-step process that includes both automated as well as manual procedures. The goal of this process was to build an accurate archive of all the *historical versions* of platform policies. That means, we always included a new instance of a document into the archive, when we *detected a change in the policy text* of the document. This is a crucial decision for the setup and implementation of the data collection process, and for the characteristics of the resulting PGA dataset. In consequence, a high number of document instances in a given year for a given document (eg. Community Guidelines of Facebook in 2017) means that in this year these policies have been changed quite frequently.

Our data collection process included four steps that we detail in the following:

- 4.1. Building URL timelines for the policy documents
- 4.2. Automated scraping of URL snapshots from the Internet Archive's Wayback Machine
- 4.3. Automated checking for changes (diff-check)
- 4.4. Manual revision of the automatically identified versions

In the case of interactive pages and multi-page documents such as Facebook's Community Guidelines after 18 March 2015 and Twitter's Community Guidelines after 30 May 2013, the automated steps of this process could not be implemented. In these cases we therefore opted for a different process which involves the manual collection of changed versions and the manual compilation of the different subsites into one document. This process is further detailed in section 4.5 of this chapter.

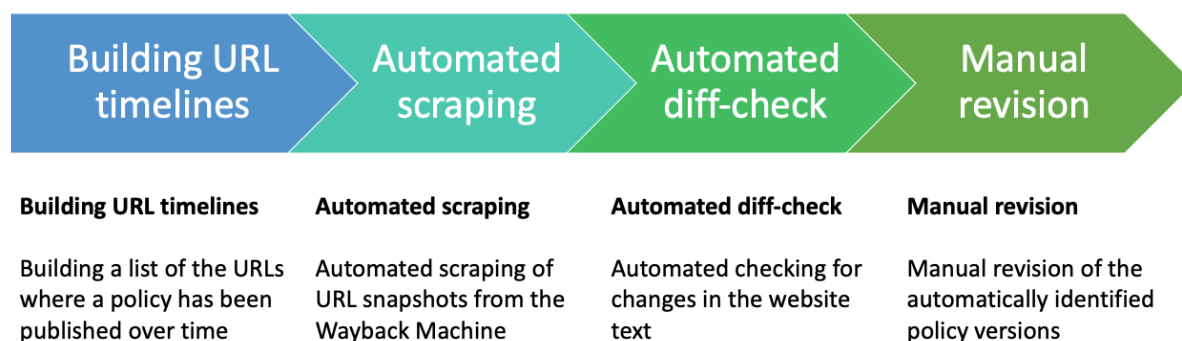


Figure 1: Steps of the Data Collection and Data Cleaning Process

4.1 Building URL timelines

The first step in the data collection process consisted of mapping how the URLs of a given policy changed over time. While this may sound like a straightforward and easy task, it is actually a form of platform archaeology or historiography (Helmond & Vlist, 2019) and consists of putting together the puzzle of how a platform's site architecture evolved over time and how policies were renamed and relocated. With the help of the Internet Archive's Wayback Machine, we therefore put together a timeline for each policy which contains the different URLs at which it has been located. The resulting table (cf. Table 2) contains the date of when a URL was introduced as well as the date when it was replaced by a new URL. These dates serve as the starting and end point for the automated scraping process.

The Wayback Machine is an online tool by the non-profit organisation "The Internet Archive" (2023a). Its aim is to create an archive of the web and to make it accessible for researchers, historians and scholars (The Internet Archive, 2023c). As of now, it contains historical records of over 808 billion public webpages (The Internet Archive, 2023b). The records of websites, so called "snapshots" or "captures" come from different sources who crawl the web and contribute to the overall collection (The Internet Archive, 2023c). They can be accessed via an online interface which allows users to browse through the history of a given URL and directly access and navigate all saved historical versions. The Wayback Machine also offers an API which we used for the automated scraping (4.2.) of the PGA process (The Internet Archive, 2013).

In order to identify the past URLs of a given platform policy we developed the following procedure: Starting from the current URL of the policy, we used the Wayback Machine to trace back when it was first introduced by checking for the earliest available snapshot of the URL. We took this date as the inception date of the URL. We then used the Wayback Machine's records of the overarching platform site to find out where the policy was located prior to the inception of its current URL. At times, there was an overlap between two successive URLs and they existed in parallel for some time. In this case, we used the inception date of the newer URL as the breaking point because the record of the newer

URL was often more comprehensive and complete. As the end point of the older URL, we used the date of the latest available Wayback Machine record before the inception date of the more recent URL.

Platform	Document type	URL	From	To
Facebook	Terms of Service	http://www.thefacebook.com/terms.php	2004/06/11	2005/08/31
Facebook	Terms of Service	https://www.facebook.com/terms.php	2005/11/26	2011/11/29
Facebook	Terms of Service	https://www.facebook.com/legal/terms	2011/11/30	2021/12/31
Facebook	Community G.	http://www.facebook.com/codeofconduct.php	2007/08/30	2010/08/20
Facebook	Community G.	https://www.facebook.com/communitystandards/	2011/01/27	2015/03/14
Facebook	Community G.	Manual compilation	2015/03/18	2021/12/31
Facebook	Privacy Policy	http://www.thefacebook.com/policy.php	2004/06/13	2005/06/24
Facebook	Privacy Policy	http://www.facebook.com/policy.php	2005/08/09	2011/09/06
Facebook	Privacy Policy	http://www.facebook.com/full_data_use_policy	2011/09/15	2015/01/20
Facebook	Privacy Policy	http://www.facebook.com/policy.php	2015/01/30	2021/12/31
YouTube	Terms of Service	http://www.youtube.com/terms.php	2005/04/28	2005/11/02
YouTube	Terms of Service	http://youtube.com/t/terms	2005/12/17	2021/12/31
YouTube	Privacy Policy	http://www.youtube.com/privacy.php	2005/04/28	2005/11/02
YouTube	Privacy Policy	http://youtube.com/t/privacy	2006/04/09	2012/02/15
YouTube	Privacy Policy	http://www.google.com/intl/en/policies/privacy/	2012/07/27	2018/03/22
YouTube	Privacy Policy	https://policies.google.com/privacy?hl=en	2018/03/23	2021/12/31
Twitter	Terms of Service	http://twitter.com/tos	2006/11/09	2017/06/18
Twitter	Terms of Service	https://twitter.com/en/tos	2017/06/19	2021/12/31
Twitter	Community G.	Manual compilation	2009/01/18	2021/12/31
Twitter	Privacy Policy	https://twitter.com/privacy	2008/11/19	2017/06/13
Twitter	Privacy Policy	https://twitter.com/en/privacy	2017/06/14	2021/12/31
Instagram	Terms of Service	http://instagram.com/legal/terms	2011/04/18	2012/04/20
Instagram	Terms of Service	http://instagram.com/about/legal/terms/	2012/04/30	2014/10/30
Instagram	Terms of Service	https://help.instagram.com/478745558852511	2014/10/31	2018/04/22
Instagram	Terms of Service	https://help.instagram.com/581066165581870	2018/04/23	2021/12/31
Instagram	Community G.	http://help.instagram.com/customer/portal/articles/262387-community-guidelines	2012/05/03	2012/10/06
Instagram	Community G	https://help.instagram.com/477434105621119	2012/12/30	2021/12/31
Instagram	Privacy Policy	http://instagram.com/legal/privacy	2011/07/18	2012/12/08
Instagram	Privacy Policy	http://instagram.com/about/legal/privacy/updated/	2012/12/18	2013/03/12
Instagram	Privacy Policy	https://help.instagram.com/155833707900388	2013/03/17	2018/04/22
Instagram	Privacy Policy	https://help.instagram.com/519522125107875	2018/04/23	2021/12/31

Table 2: URL List of collected platform policies

As an example: The current URL of Twitter's English Terms of Service is <https://twitter.com/en/tos>. The earliest intact snapshot of this URL on the Wayback Machine is from 19 June 2017. This date therefore marks the inception date of this URL and the starting point for the automated crawling process. To find out where Twitter's Terms of Service were located before the inception of this URL, we opened the latest intact snapshot of the complete site (<https://twitter.com/>) before 19 June 2017 which is from 18 June 2017. Navigating this snapshot and clicking on the "Terms" button on the bottom of the page leads to the older URL of Twitter's Terms of Service: <https://twitter.com/tos>. Using the Wayback Machine's calendrical overview of this older URL shows that its first intact record is from 9 November 2006. Since the newer URL exists from 19 June 2017 onwards we subsequently check for the latest available intact record of the older URL before this date which is 18 June 2017. The older URL is hence crawled from 9 November 2006 to 18 June 2017, while the newer URL (<https://twitter.com/en/tos>) is crawled from 19 June 2017 until 31 December 2021.

During this process, we also came across the fact that policies were sometimes renamed, for instance from 'Terms of Use' to 'Terms of Service' or from 'Privacy Policy' to 'Data Policy'. A timeline of these name changes is given in the following three tables. For the naming of the files in the dataset we chose to go for the uniform labels of *Terms of Service*, *Community Guidelines*, and *Privacy Policy*.

Platform	Policy title	Timeframe
Facebook	Terms of Use	2004/06/11-2009/02/04
	Statement of Rights and Responsibilities	2009/06/03-2018/05/26
	Terms of Service	2019/08/03-2021/10/30
YouTube	Terms of Use	2005/04/28-2007/09/26
	Terms of Service	2009/03/20-2021/11/27
Twitter	Terms of Service	2006/11/09-2021/10/28
Instagram	Terms of Use	2011/04/18-2021/10/28

Table 3: Name development of Terms' of Service

Platform	Policy title	Timeframe
Facebook	Privacy Policy	2004/06/13-2006/10/31
	Data Use Policy	2011/09/15-2013/11/19
	Data Policy	2015/02/03-2021/11/14
YouTube	Privacy Policy	2005/04/28-2005/10/31
	YouTube Privacy Notice	2008/11/20-2010/12/17
	Privacy Policy	2012/07/27-2017/12/18
	Google Privacy Policy	2018/06/01-2021/07/03
Twitter	Twitter Privacy Policy	2008/11/19-2021/08/19
Instagram	Privacy Policy (Burbn)	2011/07/18-2013/03/17
	Data Policy	2018/07/31-2021/01/13

Table 4: Name development of Privacy Policies

Platform	Policy title	Timeframe
Facebook	Content Code of Conduct	2007/08/30
	Community Standards	2015/03/18-2021/12/23
Twitter	The Twitter Rules	2009/01/18-2021/12/21

Instagram	Community Guidelines	2012/05/03-2021/10/28
-----------	----------------------	-----------------------

Table 5: Name development of Community Guidelines

4.2 Automated scraping from the Wayback Machine

Based on this list of URLs we initiated an automated scraping of the Wayback Machine. Given one specific URL we first queried all snapshots of this URL which were stored by the Wayback Machine. For example, the Privacy Policy of Twitter was accessible at <https://twitter.com/en/privacy> from 14 June 2017 until the end of 2021. The query for the history of this URL at the Wayback Machine can be sent via the following http request:

```
(GET)
http://web.archive.org/cdx/search/cdx?url=https://twitter.com/en/privacy&output=json&collapse=digest&filter=statuscode:200
```

This request uses the `digest` field and the `statuscode` field (a value of 200 indicates a valid response from the server) to filter out duplicates and error-pages. This already marks the first automated step of document filtering. The value of the `digest` field is the output of a hash function which encodes the document. Even a slight change to the recorded snapshot will result in a very different digest value. Thus, we can use the digest values to filter out exactly matching snapshots before downloading the documents.

The Wayback Machine then returns a table with the fields:

```
["urlkey", "timestamp", "original", "mimetype", "statuscode", "digest", "length"]
```

This table may contain thousands of entries. Each entry represents a snapshot of the URL at a certain time. However, due to the `digest` and `statuscode` filters, the table should contain no digest twice and no entry with a status code which is not 200. Furthermore, one specific snapshot can be accessed through the unique pair of the URL and the “timestamp” field of the table.⁷ In our case this might be:

```
http://web.archive.org/web/20180313021009/https://twitter.com/en/privacy
```

However, even with those two filters in place, most of the downloaded documents are still equal in the terms of the policy text. This is due to the `digest` being very sensitive to changes that have nothing to do with the text of the document. Just a different ad picture on the website, a change of colours or the text font will change the `digest` and thus yield a new version, even if the difference is not human-perceivable. Thus, the resulting documents of these steps do not yet represent the corpus of all historical versions of platform policies and need additional steps of data cleaning.

In order to archive a manageable dataset for further processes, we only archived one snapshot of a policy per calendar week even if there were more than one record available. We downloaded those documents in PDF format using the `pdftkit` library (<https://pypi.org/project/pdftkit/>). The resulting dataset thus consists of PDF snapshots of the full webpages containing the policies. (<https://github.com/PlatformGovernanceArchive/public-pga-workbench>)

4.3 Automated checking for changes

As a next step we needed to isolate the policy text within the downloaded snapshots and then check for changes in this text. Detecting changes in documents can be achieved by commercial tools such as *Draftable* (<https://www.draftable.com/>). But we wanted to keep this procedure open, reproducible, and cost-effective and hence prioritised open-source libraries. This turned out to be a challenging task

⁷ The web archive provides an API with diverse options of queries. For more information, visit their documentation: <https://github.com/internetarchive/wayback/blob/master/wayback-cdx-server/README.md>).

for the following reasons.

First, using the HTML code directly was too noisy due to obfuscation tools and technology changes (e.g. PHP to HTML5) in the webpages of the platforms. Secondly, extracting only the text from the PDF document would have deleted some information that is important for the context such as headlines and paragraphs. Furthermore, we wanted to display the diff-checks (the visual display of changes in the text) for better usability for other researchers.

Isolating the Text

In order to resolve both issues, we normalised all documents by converting the PDF files into HTML files using the *PyMuPDF* library (<https://pymupdf.readthedocs.io/en/latest/module.html>). Then we converted the HTML files into Markdown files, using the *Markdownify* library (<https://pypi.org/project/markdownify/>). This conversion process yields formatted text-only documents which can then be compared to each other. During this process, we deleted all non-English versions and empty documents which are sometimes part of Wayback Machine's records.

Automated Process for Identifying Changes

We then used the *diffliib* library (<https://docs.python.org/3/library/diffliib.html>) to detect changes between documents. It compares texts on a character basis. All documents that contained no change in the website text, including only spacing changes and only positional changes (line breaks), were removed from the corpus because these do not constitute new policy versions according to our definition of a new version which is laid out below.

4.4. Manual revision of the automatically identified versions

The automated filtering process narrowed down the number of possible policy versions considerably, but still a significant amount of data noise remained. This is due to the fact that the automated system 1) registered changes in all elements of the documents such as the header, footer and sidebar and not just the policy text and 2) registered changes in each non-whitespace character of the text which also included changes such as mere typo corrections which we did not consider relevant enough to constitute a *new version* of a policy within the PGA v1 corpus.

For this reason, we additionally conducted a thorough manual revision of the results of the automated process which only focused on the policy text. In order to do so, we first had to define what constituted the *policy text* as well as a *new policy version* from our perspective. This entails defining what constitutes a *relevant change*. The goal of this definition was to record all substantive changes to the policies that might have practical, legal or normative repercussions while discarding meaningless changes. The documents were filtered based on the following definition of a new policy version:

We detect a new version of a policy if at least one word in the policy text changes. The policy text does not include other elements such as the header, footer and sidebar on a site. This definition excludes mere changes to the punctuation, formatting (e.g. spacing, change of the type of quotation marks) and capitalization of words in the policy text. We furthermore exclude pure typo corrections, changes to the contact information of platforms (such as a change of the company address or company email) and mere changes in company site names (such as if a "Help page" is renamed as "Support page") from constituting a new version. Since it is the policy text which we looked at, changes of images and other graphics on the site as well as changes to the hyperlinks behind the text by themselves also do not suffice for a new policy version.

Table 6: Definition of a new policy version

Based on this definition of what constitutes a new version of a policy we were able to acquire a uniformly identified data set by manually revising the output of the automated scraping and filtering process.

The last step of data cleaning consisted of removing all elements from the automatically downloaded

PDF documents which are not part of the policy text. In order to achieve this, we manually went back to the Wayback Machine record of each unique new policy version and produced a clean PDF by saving only the policy text as a PDF. In the rare case that a record included both a new and an older version of a policy which appeared on top of each other, we only saved the text of the newer version of the policy.

It is important to note that the timestamps of the documents in the archive denote the point in time when they were registered by the automated system and that this does not necessarily match the exact date of when they were published or became effective. As the scraping was limited to a maximum of one new version per week, there are often a few days in between the actual publishing date and the date of registration by the system.




Figure 2: Document filtering process of the Platform Governance Archive

4.5. Collecting interactive pages and Multi-page documents

Two special cases of how policies were displayed on the platforms' sites complicated the collection: Interactive pages and multi-page documents. *Interactive pages* contain elements such as drop-down sections/accordions which require some form of interaction with the site in order to display the whole policy text. The content of these policies could not be automatically scraped from the Wayback Machine's snapshots. This has been the case for YouTube's Community Guidelines after 12 November 2008 (YouTube, 2008a) and Facebook's Community Guidelines after 18 March 2015 (Facebook, 2015).

Multi-page documents refer to the case in which a policy is not located at one URL but instead spread across different web pages, making them uncollectible for the automated setup of the system. This has been the case for Twitter's Community Guidelines after 1 March 2010 (Twitter, 2010) and Facebook's Community Guidelines after 24 April 2018 (Facebook, 2018). After the policies were transformed into an interactive or a multi-page format we therefore opted for a manual collection process which is detailed in the following sections.

The fact that platform policies can be spread across different URLs and that especially the Community Guidelines seem prone to evolve into multi-page documents over time raises an important ontological and definitional question which is what a platform's Community Guidelines actually consist of. Are the Community Guidelines of a platform what the company defines as such or do they actually encompass all of the platforms' rules that regulate the "community" of their users in some way? This would mean that rules which are spelled out in parts of a site which are not part of the officially designated "Community Guidelines", such as on help pages, would also form part of a platform's Community Guidelines. For reasons of feasibility and practicability, we opted for taking the platform's

own definition of their "Community Guidelines" as the reference point for our collection.

The case of Facebook's Community Guidelines can help to illuminate these conceptual questions. Initially called "Content Code of Conduct" and later renamed as "Facebook Community Standards", they have historically evolved from a document that was completely displayed on one URL into first an interactive document with drop-down sections/accordions and then a multi-page document which spreads across many URLs. But although the document is now spread across different URLs, its historical evolution makes it clear that the multi-page format should not be seen as a splitting up of the Community Guidelines into different sub policies but much rather the contemporary form of displaying the document and making it easier to navigate for users.

For the collection of Facebook's Community Guidelines after 18 March 2015, we hence opted for collecting the content of the different subsections and compiling it back into one document. However, the sheer number of different URLs and holes in the records of the Wayback Machine made it impossible to collect them through the aforementioned process. We hence manually collected the text of the different sections from the Wayback Machine's snapshots and identified new versions by comparing the text of the different compiled versions with an automated diffchecker based on the PGA's definition of new policy versions.

For the versions after 30 April 2018, we relied on the public archive that Facebook has been offering which shows changes that were made to its Community Guidelines after this date (Meta, 2023a). Whenever a change to a section was listed in Facebook's "change log", we hence registered a new version of the Community Guidelines for the corpus. Further details on the specificities of the collection process of Facebook's Community Guidelines and the format of the resulting documents are given in section 5.3.1. of the following chapter.

The ontological question of what the Community Guidelines encompass was harder to answer in the case of Twitter. Twitter's Community Guidelines have ever since been called "The Twitter Rules" and began to include links to subpages as early as 1 March 2010. Since 6 December 2017 Twitter has also maintained a "Rules and policies" page which, at the time, linked to 40 sub pages (Twitter, 2017b) and currently gives an overview of 84 subsites while in the current version of "The Twitter Rules" only 18 selected sub policies are linked (Twitter, 2023b).

We decided to consider the "Twitter Rules" page as the departure point for our collection because it is what Twitter has historically considered its central set of rules for the community and is the document that all the other rules emerged from. This decision was supported by the fact that Twitter's "Rules and Policies" page includes many usage guidance and information pages such as "Updates to our Terms of Service and Privacy Policy" which can rather be classified as help pages than as rules or policies for the community. We hence decided to collect "The Twitter Rules" as well as the content of the first sublevel of the policies that were linked on this page. The details of this manual collection process are laid out in section 5.3.3. of the following chapter.

Finally, YouTube's Community Guidelines were not collected as part of this dataset because the policy was especially fragmented and because the Wayback Machine's records of its subpages were very incomplete.

5 The PGA v1 Corpus

The resulting final corpus of the PGA is a dataset consisting of 354 policy documents with a total of 6,036 pages, representing all historical versions of the selected platform policies by Facebook, Instagram, Twitter, and YouTube from the inception of these policies until the end of 2021. The following sections describe how we store and provide this dataset (5.1), and detail the specificities of each platform and policy type (5.2).

5.1 Project Website and Project Elements

The project website and archive can be accessed via <https://platformgovernancearchive.org/>. There

you find all relevant information on the project, team members, links to downloading datasets, as well as finding an interactive interface to explore the archive and the changes between policy versions through a timeline visualisation (cf. Figure 2)⁸.

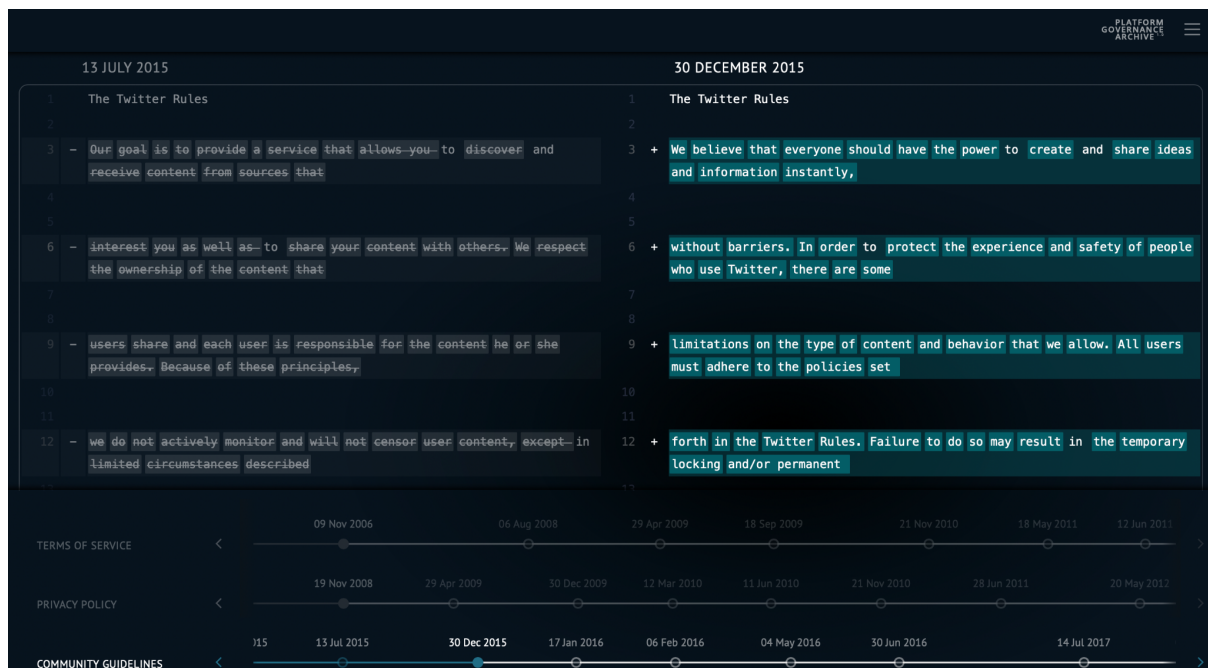


Figure 3: Visual interface of the Platform Governance Archive

The website also provides a download option to bulk download the full PGA v1 corpus. For more detailed access to our dataset, the website links to our repository at Github: <https://github.com/PlatformGovernanceArchive>.

5.2 Data Repositories: Folder Structure and Naming Conventions in the Corpus

The Github repositories of the Platform Governance Archive provide detailed access to the PGA v1 corpus as well as to research data and instruments that we used in the process, including the URL list, scripts and the datasets before the data cleaning processes (cf. 4.3).

The Platform Governance Archive v1 consists of two repositories:

/pga-corpus (the corpus of all identified policy versions in PGA v1)

/pga-workbench (providing tools, and data that we used in the research process)

The generic naming convention for files is that filename bears the timestamp of the documents. This timestamp denotes the point in time when they were registered by the system. This does not necessarily match the exact date of when they were published or became effective. As the scraping was limited to a maximum of one new version per week, there are often a few days in between the actual publishing date and the date of registration by the system. This means that a document can mention a date of effectiveness that comes before the date in the name of the file.

5.2.1 /pga-corpus

The [/pga-corpus repository](#) is structured as follows:

⁸ This visualisation was designed and implemented by team member and designer Larrisa Wunderlich. Portfolio of Larissa Wunderlich: <https://www.larissawunderlich.de/>.

Folder	Description	Format	Units	Potential Use Cases
pga-corpus/Versions	Final dataset containing all identified historical policy versions of the PGA v1	PDF (individual & merged), HTML & Markdown	1,071	Analysis of the historical evolution of platform policies; analysis of the evolution of specific policy sections/provisions

In addition to the individual versions, we also provide an overall PDF file for each policy which contains all historical versions of a policy in one file.

5.2.1 /pga-workbench

In [/pga-workbench](#) we provide the following data and tools:

Folder	Description	Format	Units	Potential Use Cases
pga-workbench/URL List	List of all the URLs at which the policies were stored over time which served as the basis for scraping the WBM	CSV	1	Reproducing our data collection process; investigating URL changes
pga-workbench/Snapshots	All policy snapshots that were scraped from the WBM as part of the first data collection step (4.2)	PDF, HTML & Markdown	15,039	Reproducing our data cleaning process, Applying a different document filtering process; investigating the visual appearance of past platform policies
pga-workbench/Diff check	Comparison view of the automated diff-check	HTML & Markdown	9,636	Reproducing our data cleaning process, Applying a different document filtering process
pga-workbench/Scripts	Scripts that were used for the automated scraping (miner), the creation of document types & automated diffchecking (convert versions) and for the creation of the plots (plots) in chapter 6	IPYNB	3	Reproducing and further understanding our data collection process
pga-workbench/Plots	Heatmaps and graphs about the development of the change frequency and character count of the overall PGA v1 corpus, as well as by platform & policy type	SVG & PNG	35	Analysis and visualisation of the overall development of the PGA corpus; identification of general trends; comparative analysis between policy types/platforms

5.3 Characteristics and Details of the Platforms and Policies

5.3.1 Facebook

For Facebook, the dataset contains 154 separate policy versions with a total of 2,862 pages. The site, initially called “Thefacebook”, was launched on 4 February 2004 (Rothman, 2015). The earliest snapshot of the site on the Wayback Machine is from 12 February 2004 (Thefacebook, 2004a).

Facebook’s Terms of Service were first recorded by the Wayback Machine on 11 June 2004 (Thefacebook, 2004b). Initially titled “Terms of Use”, they were renamed to “Statement of Rights and Responsibilities” on 3 June 2009 (Facebook, 2009). Since 3 August 2019 they have been called “Terms of Service” (Facebook, 2019).

Since 30 June 2005, the Terms of Service have included the date on which they became effective

(Facebook, 2005a). Over the years, they furthermore began to include specific regional provisions - such as a section on the "Network Enforcement Act ("NetzDG")" which is visible when accessed from a German IP.

For the PGA, only the English US/International version of the Terms of Service has been collected. This version has, at times, included provisions which explicitly refer to legislation in the United States such as references to laws of the State of Massachusetts and the California Civil Code which was introduced with the first version (Facebook, 2004).

Policy type	First recorded	Number of versions	Details
Terms of Service	2004/06/11	38 versions, 239 pages	<ul style="list-style-type: none"> - Renamed from "Terms of Use" (2004/06/11-2009/02/04) to "Statement of Rights and Responsibilities" (2009/06/03-2018/05/26) to "Terms of Service" (2019/08/03-end of 2021) - Date of effectiveness since 2005/06/30 - Regionalized provisions
Privacy Policy	2004/06/13	43 versions, 301 pages	<ul style="list-style-type: none"> - Renamed from "Privacy Policy" (2004/06/08-2010/12/24) to "Data Use Policy" (2011/09/15-2013/11/19) to "Data Policy" (2015/02/03-end of 2021) - Date of effectiveness since 2005/08/09 - Regionalized provisions
Community Guidelines	2007/08/30	73 versions, 2,322 pages	<ul style="list-style-type: none"> - Renamed from "Content Code of Conduct" (2007/05/24) to "Community Standards" (2011/01/27-today) - Manual collection after 2015/03/18 (Interactive page since 2015/03/18, Multi-page document since 2018/04/24) - Versions after 2018/04/30 based on company change log

Facebook's Privacy Policy was first recorded by the Wayback Machine on 13 June 2006 (Thefacebook, 2004c). It was renamed to "Data Use Policy" with the version from 15 September 2011 (Facebook, 2011b). With the version from 3 February 2015, the document title was shortened to "Data Policy".

Since the version from 9 August 2005, the Privacy Policy has included a date of effectiveness (Facebook, 2005b). Furthermore, it has also begun to include regionalized provisions. For instance, current versions which are accessed from a European Union IP include a section on the General Data Protection Regulation (GDPR). As part of the PGA, the US/International version has been collected.

Facebook's Community Guidelines were first recorded by the Wayback Machine on 30 August 2007 (Facebook, 2007). At the time of their inception, they were called "Content Code of Conduct" (Facebook, 2007). Since the version from 27 January 2011, they have been referred to by the platform as "Facebook Community Standards" (Facebook, 2011a).

As laid out in section 4.5., Facebook's Community Guidelines were transformed into an interactive page with drop-down sections on 18 March 2015 (Facebook, 2015) and into a multi-page document on 24 April 2018 (Facebook, 2018). As these page layouts were uncollectable for our automated process, the Community Guidelines have since been manually collected.

The versions between 18 March 2015 and 30 April 2018 were collected and identified by manually

compiling the text of the different drop-down sections from the Wayback Machine records into one document and then comparing them with an automated diffchecker based on the PGA's definition of a new version.

The versions after 30 April 2018 were put together based on the public "change log" that Facebook has been offering of changes that were made to its Community Guidelines after April 2018 (Meta, 2023a). In line with the change definition of the PGA, we registered a new version of the overarching Community Guidelines document, if the text of at least one of its sections changed.

The versions after 30 April 2018 hence carry the dates of change that are visible within Facebook's interactive archive and their content was manually collected from this site and compiled into text documents in a unified formatting. Facebook's visualisation of changes which shows additions as marked in green and removals as strikethrough text were removed as part of this process to keep the documents comparable to the rest of the archive.

The explanatory footnotes which Facebook offers about the context of specific changes as part of its archive and which sometimes include a detailed explanation of the company's reasoning behind a change, have been included in the manually collected versions. Although this marks an exception to the rest of the archive, we did not want to remove this additional context. The footnotes always begin with the phrase "Note from Facebook's records" and carry the date on which they were made.

Additional explanatory footnotes were added by us whenever there was an unclarity in the manual collection process and begin with "Note by the Platform Governance Archive".

5.3.2 YouTube

The dataset for YouTube consists of 43 policy versions with a total of 330 pages. YouTube was launched to the public in May 2005 (YouTube, 2011) and the main site was first recorded by the Wayback Machine on 5 May 2005 (YouTube, 2005d). Subsites such as the first version of the Terms of Service, however, were already recorded earlier, in this case on 28 April 2005 (YouTube, 2005c).

Initially called "Terms of Use", they were renamed as "Terms of Service" with the version from 20 March 2009 (YouTube, 2009). While the first version of the policy included a date of effectiveness (YouTube, 2005c), not all following versions have carried such a date since then.

Regionalized versions of the Terms of Service have existed at least since 2010, when the versions between the United States and the United Kingdom began to diverge. At least since the version from 2 June 2018 (YouTube, 2018), there have been different versions for the European Economic Area and countries outside of this area, including the United States. For the PGA, the US/International version has been collected.

Policy type	First recorded	Number of versions	Details
Terms of Service	2005/04/28	14 versions, 79 pages	<ul style="list-style-type: none"> - Renamed from "Terms of Use" (2005/04/28-2007/09/26) to "Terms of Service" (2009/03/20-end of 2021) - Date of effectiveness missing in some versions - Regionalized versions at least since 2010
Privacy Policy	2005/04/28	29 versions, 251 pages	<ul style="list-style-type: none"> - Renamed from "Privacy Policy" (2005/04/28-2005/10/31) to "YouTube Privacy Notice (2008/11/20-2010/12/17) to "Privacy Policy (2012/07/27-2017/12/18) to "Google Privacy Policy (2012/07/27-end of 2021) - Date of effectiveness/last revision since first version

Community Guidelines	//	//	- Not collected
----------------------	----	----	-----------------

YouTube's Privacy Policy was first recorded on the Wayback Machine on 28 April 2005 (YouTube, 2005a). Initially called "Privacy Policy" (YouTube, 2005b), the title of the document was changed to "YouTube Privacy Notice" with the version from 20 November 2008 (YouTube, 2008b) until it was again named "Privacy Policy" from 27 July 2012. Since the version from 1 June 2018, the document has been titled "Google Privacy Policy" (Google, 2018).

The policy had already begun to reference Google's Privacy Policy in the version from 20 November 2008 (YouTube, 2008b) and continued to do so up until the version from 27 July 2012 (Google, 2012). With this version, it seems that the company ceased to maintain two separate Privacy Policies and just offered the Google Privacy Policy for YouTube users. The Privacy Policy has included a date of effectiveness or of the last revision since its first version.

YouTube's Community Guidelines have not been collected as part of the Platform Governance Archive. This is due to the fact that they began to include interactive elements as early as on 12 November 2008 (YouTube, 2008a) which could not be collected automatically.

5.3.3 Twitter

For Twitter, the PGA v1 corpus contains 126 policy documents with a total of 2,708 pages. Initially called "Twtrr", the service was publicly launched on 15 July 2006 (Arrington, 2006). The first snapshot of the site on the Wayback Machine is from 30 September 2006 (Twitter, 2006a).

Twitter's Terms of Service were first recorded by the Wayback Machine on 9 November 2006 (Twitter, 2006b). They have included an effective date since the version from 18 September 2009 (Twitter, 2009d).

At least since 17 September 2009, Twitter has publicly displayed previous versions of its Terms of Service on its site (Twitter, 2009c). However, this corporate archive only lists 16 versions of the Terms of Service until the end of 2021 (Twitter, 2023a) while the PGA has identified 21 new versions in this period.

With the version from 5 October 2016 (Twitter, 2016), Twitter introduced two separate versions of its Terms of Service for users who live in the United States and for those who live "in the European Union or otherwise outside the United States". The company has been maintaining these two regionalized versions since then. As they have been displayed on top of each other on the same URL, both regional versions have been recorded as part of the PGA.

With the version from 6 January 2020, Twitter extended the scope of its EU/Non-US Terms of Service to the whole "European Economic Area" (Twitter, 2020a) and with the version from 19 June 2020, the United Kingdom was added to its area of applicability (Twitter, 2020b).

Policy type	First recorded	Number of versions	Details
Terms of Service	2006/11/09	21 versions, 139 pages	<ul style="list-style-type: none"> - Date of effectiveness since 2009/09/18 - Company archive of previous versions since 2009/09/17 - Regionalized versions for the US and EU/outside of the US since 2016/10/05 - Version for EU/EFTA states and outside EU/EFTA (including US) since 2020/01/06
Privacy Policy	2008/11/19	21 versions, 129 pages	<ul style="list-style-type: none"> - Date of effectiveness since first version - Company archive of previous versions since

			2017/06/16
Community Guidelines	2009/01/18	84 versions, 2,386 pages	<ul style="list-style-type: none"> - Multi-page document since 2010/3/01 - Manual collection since first version - Sub policies sometimes include a date of last revision

Twitter's Privacy Policy was first recorded by the Wayback Machine on 19 November 2008 (Twitter, 2008a). It has been called "Twitter Privacy Policy" ever since and has included a date of effectiveness from its very first version (Twitter, 2008b).

An overview of previous versions of the Privacy Policy has existed on Twitter's site at least since 16 June 2017 (Twitter, 2017a). However, the company's archive only lists 17 previous versions of the Privacy Policy until the end of 2021 (Twitter, 2017a), while the PGA has identified 21 versions in this period.

Twitter's Community Guidelines have been referred to by the platform as "The Twitter Rules" since their first version, which was recorded by the Wayback Machine on 18 January 2009 (Twitter, 2009b). They were transformed into a multi-page document as early as 1 March 2010 (Twitter, 2010) when the first two links to policies outside of the "Twitter Rules" appeared in the policy text. With the version from 30 May 2013 (Twitter, 2013), many more URLs appeared in the document which linked to policies on topics such as "Impersonation", "Violence and Threats" and "Copyright".

As mentioned in section 4.5., we therefore opted for a manual collection of Twitter's Community Guidelines from the first version. In practice, this meant that we first had to create a timeline which denoted when a link to a subpolicy was included or removed from the overarching Twitter Rules page. In this context, it is important to note that some of the linked policies already existed before they were included in the Twitter Rules page or continued to exist after they were removed from it. An inclusion or removal of a policy into/from the "Twitter Rules" therefore does not equate the inception or cessation of this policy

We collected the content of the different subsites of the "Twitter Rules" from the Wayback Machine's snapshots of these URLs and then checked them for changes with an automated diffchecker. A new version of Twitter's Community Guidelines was hence registered when there was a change in the policy text of either the overarching Twitter Rules document or one of the linked subpolicies, or when subpages were added or removed from the master document.

The resulting compiled documents always begin with the record of the overarching "Twitter Rules" page, followed by the content of the different subpolicies. The underlying Wayback Machine records are linked in the section titles.

Each section carries its version date in brackets behind the title. Whenever the exact date of a new version could not be identified - mostly due to an incompleteness of the Wayback Machine's records - it is followed by a star symbol (*). To facilitate seeing which section in the "Twitter Rules" changed, the title of changed sections are marked in yellow.

5.3.4 Instagram

For Instagram, the corpus contains 31 policy versions with a total 136 pages. Instagram was first launched as an iOS application on 6 October 2010 (Siegler, 2010). The first snapshot of Instagram's web domain on the Wayback Machine is from 11 March 2011 (Instagram, 2011a).

The platform's Terms of Service were first recorded on 18 April 2011 (Instagram, 2011b). They have been consistently referred to as "Terms of Use" by the platform and began to include a date of effectiveness as well as a link to the previously valid version with the version from 3 February 2013 (Instagram, 2013).

At least since 15 May 2018, Instagram has offered regionalized versions of its Terms of Service. On this date, the Wayback Machine recorded a version of the Terms of Use with a specific section for

users from the European Union (Instagram, 2018a). The PGA recorded the International version of the Terms of Service which does not contain such a provision.

Instagram's Privacy Policy was first recorded on the Wayback Machine on 18 July 2011 (Instagram, 2011c). It was renamed "Data Policy" with the version 31 July 2018 (Instagram, 2018b). The policy has included a date of effectiveness since the version from 7 September 2012 (Instagram, 2012c).

Instagram's Community Guidelines were first recorded on the Wayback Machine on 3 May 2012 (Instagram, 2012a). Since its very first version the document has been referred to by the platform as Community Guidelines (Instagram, 2012b). Until now, the document has not included a date of effectiveness.

Policy type	First recorded	Number of versions	Details
Terms of Service	2011/04/18	6 versions, 37 pages	<ul style="list-style-type: none"> - Date of effectiveness since version 2013/02/03 - Regionalized versions at at least since 2018/05/15
Privacy Policy	2011/07/18	9 versions, 62 pages	<ul style="list-style-type: none"> - Renamed from Privacy Policy (2011/07/18-2013/03/17) to Data Policy (2018/07/31-2021/01/13) - Date of effectiveness since 2012/09/07
Community Guidelines	2012/05/03	16 versions, 37 pages	<ul style="list-style-type: none"> - No date of effectiveness included


6 Trends in the Evolution of Platform Policies

The following chapter presents a structural analysis of some of the general trends and patterns which are visible in the historical evolution of the PGA dataset in a time period of up to nineteen years on the document level. We do so by looking at two different dimensions: The development of the number of characters of each policy and the frequency at which they were changed. A comparative visualisation of the quantitative analysis allows us to show how the extent of the policies has grown over time, to identify periods of high growth and frequent changes and to draw comparative conclusions about the four different platforms and the different policy types.

6.1. Character count


The following graphs display how the character count of the policies in the PGA v1 corpus has changed over time. The character count is displayed by quarter, meaning that if there was more than one policy version in a quarter, the average character count for the versions in this quarter is given.

Average character count of Facebook policies (2004-2021)




The character count of Facebook’s Terms of Service grew steadily after they were introduced in 2004 until they were cut by 46% in the fourth quarter of 2008. Since then, its character count has remained relatively stable. The company’s Privacy Policy also grew steadily until the third quarter of 2012. In the fourth quarter of 2018, the length of the Privacy Policy decreased significantly, while the Community Guidelines grew in parallel. Facebook’s Community Guidelines have grown massively since their inception, increasing by 2,548 % from 4,305 characters in August 2007 to 113,990 characters at the end of 2021. The growth has picked up significantly since the beginning of 2018, with a short but significant decrease in the second quarter of 2019.

Average character count of YouTube policies (2005-2021)




YouTube’s Terms of Service grew significantly right after they were introduced, as well as in the third quarter of 2007 and the second quarter of 2010. After that, its character count has remained very stable. The Privacy Policy has grown by 130% since its inception until the end of 2021, with a period of high growth between 2008/4 and 2009/1, reductions in 2010/4 and 2017/4 and an overall pattern of continuous growth since the third quarter of 2012.

Average character count of Twitter policies (2008-2021)




Twitter's policies grew in parallel, with the Terms of Service increasing by 612%, the Privacy Policy by 262% percent and the Community Guidelines by 3,354% from their inception until the end of 2021. While the growth of the Terms of Service and Privacy Policy occurred more gradually, the Community Guidelines saw periods of rapid growth in the first quarter of 2010, the second quarter of 2013, the fourth quarter of 2017 and an extended period of high growth from the fourth quarter of 2018 until the end of 2021.

Average character count of Instagram policies (2011-2021)




Instagram's policies have shown low levels of fluctuation when compared to the other platforms. From their inception until the end of 2021, the length of the platforms' Terms of Service increased by 233%, the Privacy Policy grew by 372% and the Community Guidelines by 35% percent. The extent of the Terms of Service declined in the second quarter of 2018. This decline was followed by an increase in the Privacy Policy in the third quarter. The Privacy Policy primarily grew in two periods: Right after its inception until the fourth quarter of 2012, as well as in the third quarter of 2018. The length of the Community Guidelines has remained relatively stable since 2013, which stands out in comparison to the other platforms, although this might be due to the fact that the document has increasingly linked to Facebook's Community Guidelines.

Average character count of Terms of Service (2004-2021)




Taking a comparative look at the growth of the Terms of Service across the four platforms, it becomes clear that Twitter has the most extensive Terms of Service document. Across the four platforms, the development of the character count of the Terms of Service seems to follow a similar pattern: A rapid increase in the first years/quarters after the inception of the document which is followed by a stabilization of the character count. The growth of the character count of Twitter’s Terms of Service after 2016/4 including the steep increase in 2018/2 represents an exception to this pattern. What also stands out is that the length Facebook’s, YouTube’s and Instagram’s Community Guidelines seems to have converged to a similar level after 2018/2.

Average character count of Privacy Policies (2005-2021)



A comparative look at the development of Privacy Policies across platforms provides an interesting insight. With the exception of Facebook before the beginning of 2015, the character count across the four platforms follows an enormously similar development pattern. Facebook’s Privacy Policy before 2015 is the outlier with a steady growth and high character counts. The steep decrease at the beginning of 2015 occurs when Facebook changed its policy from “Data Use” to “Data Policy”.

Average character count of Community Guidelines (2009-2021)



A comparative look at the development of the character count of the Community Guidelines shows that Facebook’s and Twitter’s Community Guidelines have followed a similar trajectory with distinct periods of high growth turning into a pattern of continuous growth after the fourth quarter of 2017. Periods of extension by Twitter seem to precede similar periods by Facebook. The relatively steady line for Instagram is more an artefact of our data collection than an actual result, as we did not manually collect the subpages which became linked in Instagram’s Community Guidelines as we did for Facebook and Twitter.

Overall, it is important to also note the limitations of using the character count as an indicator for policy development. A decreasing character count of a policy, for instance, CAN be an indicator for a decreasing scope of a policy, i.e. the removal of specific sections but it could also just show a streamlining of the policy text through putting it into more concise language. Moreover, looking at the character count at the document level and not at the whole body of policies of a platform can also mislead because a shrinking number of characters could indicate the moving of sections or provisions to a different policy.

6.1. Change frequency


Another interesting metric to look at is how often a policy has been changed by the platforms. The following heatmaps show how many versions of a policy were identified and collected in each semester as part of the PGA v1 corpus in the timeframe from when they were first introduced until the end of 2021. The graphs visualise the frequency of changes for each specific policy and on an aggregated level by policy type and by platform.

Change frequency Facebook policies (2004-2021)



Facebook’s Terms of Service have been changed regularly since 2007. The time period when the company released most new versions was in the three years between the second half of 2008 and the first half of 2011. A similar pattern is visible for the company’s Privacy Policy which was also most frequently changed between the beginning the second half of 2008 and the second half of 2010. It has since been adapted regularly with outliers in terms of change frequency in the second semester of 2019 and the second half of 2020. The change frequency of Facebook’s Community Guidelines has picked up rapidly since the year 2018. The highest concentration of new versions occurred in the one and a half year period between 2019/2 and 2020/2 when 47 new versions were released.

Change frequency of YouTube policies (2005-2021)



YouTube’s Terms of Service were adapted four times in the first year after they were instituted. In the years after, they were only occasionally changed, although the pace of change seems to pick up since the beginning of 2020. The frequency of changes to the platform’s Privacy Policy has picked up since the second half of 2012 with a high number of changes in 2015, as well as in the one and a half years between 2016/2 and 2017/1 and in the second half of 2020.

Change frequency of Twitter policies (2006-2021)




The pace of change to Twitter’s Terms of Service has been relatively stable over the years with a small exception in the first half of 2018 when three new versions were registered. The same holds true for the company’s Privacy Policy, although we can see periods of more frequent changes between the second half of 2008 and the first half of 2011 and between the second half of 2014 and the first half of 2017. The Community Guidelines show a different pattern with the pace of change picking up significantly with the second half of 2017. Since then, they have been adapted 52 times.

Change frequency of Instagram policies (2011-2021)



The case of Instagram shows a comparatively low change frequency for all three policies with longer periods of no new policy versions. For the Terms of Service, changes occurred only between 2011 and the first half of 2013, as well as in the first half of 2018 and the second half of 2020. The change frequency of the Privacy Policy shows a similar development pattern, with changes occurring only between the second half of 2011 and the first half of 2013, as well as the first half of 2018, and between 2020 and the first half of 2021. The pattern for the Community Guidelines looks a bit different with more constant changes occurring between 2021 and the first half of 2012 and since the beginning of 2019.


Change frequency of Terms of Service (2004-2021)



A comparative perspective on the change frequency of the Terms of Service across the four platforms


shows that Facebook stands for changing the document 38 times. In the comparable time period after 2011/1, however, when all four platforms had their Terms of Service in place, this difference narrowed. In the eleven years since then Facebook released 15 new versions and Twitter 16 new versions, while Instagram and Youtube both changed their Terms of Service seven times.

Change frequency of Privacy Policies (2004-2021)




Regarding the Privacy Policy, Facebook again stands out as changing its policy most often. In the comparable timeframe after the second half of 2011, when all Privacy Policies had been introduced, Facebook (23 changes) and YouTube (24 changes) exhibit by far the highest change frequency, followed by Twitter (14 changes) and Instagram (9 changes). There were also certain semesters when all existing platforms adapted their Privacy Policies such as 2008/2, 2009/1, 2010/2 and 2020/1.

Change frequency of Community Guidelines (2007-2021)




When comparing the change frequency of the Community Guidelines across Facebook, Instagram and Twitter, we see that Facebook published by far the most versions (93) in the comparable time frame after the first half of 2012, followed by Twitter (72) and Instagram (16). Facebook and Twitter have followed a similar development, with the frequency of new versions picking up with/after the second half of 2017, while Instagram has not followed this development and has introduced new versions of its Community Guidelines much less frequently. As mentioned before, this is more an artefact of our data collection than an actual result, as we did not manually collect the subpages which became linked in Instagram’s Community Guidelines as we did for Facebook and Twitter.

Aggregated change frequency of platform policies (2007-2021)



Looking at the aggregated frequency of new versions of the three collected policy types, it shows that Facebook has changed its policies most frequently. Again, there is a similar pattern between Facebook and Twitter with the change frequency across policies picking up significantly after/with the second half of 2017. This is a strong indicator for a growing politicisation of platform policies (Katzenbach, 2021). Instagram stands out in this regard as it not changing its policies very often. YouTube could not be included in this comparison because we did not collect its Community Guidelines.

Aggregated change frequency of platform policies (without Community Guidelines)

(2007-2021)

When only looking at platform policies without the Community Guidelines, the trend is different. We do not see this increasing frequency of change 2017-18. While terms of service, privacy policies and copyright policies constitute the core legal baseline for the operation of a platform, community guidelines in contrast deal with much nuanced questions of handling content and interaction. As a consequence, these two different graphs indicate that there is indeed a politicization of platform's role of taking decisions on content starting 2017, as this manifests in increasing changes in platforms' community guidelines.

It is important to note that the fact that a new version of a policy was published does not allow us to draw conclusions about how extensive these changes were. One platform could, for instance, publish a completely overhauled version of a policy in one semester while another platform published four new versions in the same semester with only minor changes.

The metric of the change frequency thus only gives us an indicator of how often a platform has adopted a policy but not about the extent of these changes. Another thing to note is that the colour coding in the heatmaps is always relative within charts but not across charts.

7 Challenges and Limitations

This first comprehensive version of the Platform Governance Archive has, despite our best efforts, notablentable limitations. These start with the retrospective collection of platform policies put forward here. The Internet Archive is a resource of highest value to the public and to research, yet the Archive's automated crawlers do not find and archive all versions of all websites that have ever existed. Indeed, just a fraction of these versions is recorded, and it is not possible to know with precision how much of the Internet is or is not accessible via the WayBack Machine. Furthermore, some of the pages that appear to be recorded by the Archive are not necessarily accessible via the system's interface. It was not uncommon, during our data collection, to receive an error message when trying to download a given snapshot of a website, for instance. Furthermore the density of Wayback Machine's records varies, meaning that over some time periods policies were archived many times a day while in other time periods only a couple of times per month. It is unclear if this is an artefact of irregular scraping or an indicator that pages have not changed.

Another kind of difficulty concerns the nature of the documents we were trying to collect. Platforms are increasingly documenting their policies on nested interactive pages that yield automated data collection via the Wayback Machine's API is increasingly difficult. It is also worth noting that platforms sometimes used non-textual content to explain their policies, such as videos. Seemingly designed to facilitate users' understanding of platforms' rules, these videos do not appear to be recorded by the Internet Archive and thus are not part of our collection.

Looking more closely, each platform is designed in its own way and has confronted us with specific challenges. While it is impossible for us to claim that our archive contains all versions of all policies ever published by Facebook, Twitter, and Instagram, we are quite confident in our collection of their policies. YouTube, on the other hand, presented a much harder case. Apparently due to their websites' interactive features, several URLs of YouTube's Community Guidelines do not appear to be recorded in the Archive. Furthermore, this platform has made extensive use of videos – which, as mentioned before, are not archived either.

In terms of scope, this archive only includes English language, global policies. While platforms have been remarkably late and reluctant in publishing versions in different languages and particularly with regional differences. Yet, as national and regional regulation is addressing platform governance more and more, platforms need to include rules and procedures that are specific to certain countries (e.g., Germany) and regions (e.g., EU). Language versions and regional versions of platforms' policies are not yet part of the PGA.

8 Conclusion and Roadmap

The Platform Governance Archive aims to be an accessible and sustainable resource for the research community and beyond. This data paper has documented the PGA's rationales, methods and data in its first iteration – so that researchers, journalists and other interested parties can build on this data to run their own investigations and analyses. Social media platforms have come a long way, but they are still very much in formation. In consequence, there will be continued need for research and investigations into how platforms operate, order and govern our communication.

While we publish this data paper, we are already working on the next iteration of the archive. In collaboration with [Open Terms Archive](#), a service to monitor policies across major global websites, we have established a new highly-automated data collection process, that will secure the continuous collection and update of the archive. In addition, we are expanding the scope of the archive to a wider set of platforms. These new sets of data are already available at the Github pages of our collaboration, and will be continuously updated. While Open Terms Archive is responsible for the scraping engine, we are curating the collection process by identifying relevant documents and by monitoring and cleaning the data collection process for the set of social media platforms. Another extension on our roadmap is the inclusion of other language versions as well as regional differences in the policies.

We hope that the Platform Governance Archive will become a sustainable and key resource for the emerging community of platform governance researchers and experts across the globe, and invite everyone interested to use our data as well as to reach out to us.

Funding Note

The research and data collection part has substantially benefited from funding by the European Union's Horizon 2020 research and innovation programme as part of the ReCreating Europe consortium (Grant Agreement Number 870626). The Open Science Fellowship offered by Wikimedia Deutschland to João C. Magalhães was vital for the launch of the PGA by enabling concept building and data cleaning. The Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin, and the University of Bremen have provided institutional support for the project.

Conflicts of Interest

We disclose that one of the authors, Christian Katzenbach, is a member of the expert board "Actor and Behaviour Policies" at Facebook/Meta since. He receives an annual allowance of 6.000US\$ for this engagement.

References

- Amos, R., Acar, G., Lucherini, E., Kshirsagar, M., Narayanan, A., & Mayer, J. (2020). *Privacy Policies over Time: Curation and Analysis of a Million-Document Dataset*.
<https://doi.org/10.48550/ARXIV.2008.09159>
- Arrington, M. (2006, July 16). Odeo Releases Twtr. *TechCrunch*.
<https://techcrunch.com/2006/07/15/is-twtr-interesting/>

- Barrett, B., & Kreiss, D. (2019). Platform transience: Changes in Facebook's policies, procedures, and affordances in global electoral politics. *Internet Policy Review*, 8(4).
<https://policyreview.info/articles/analysis/platform-transience-changes-facebooks-policies-procedures-and-affordances-global>
- Bipartisan Policy Center. (2022, August 8). *Technology Platforms Election Database*.
<https://bipartisanpolicy.org/blog/tech-platforms-election-database/>
- boyd, danah. (2010). Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications. In *A Networked Self*. Routledge.
- Buckley, N., & Schafer, J. S. (2022). 'Censorship-free' platforms: Evaluating content moderation policies and practices of alternative social media. *For(e)Dialogue*, Vol 4, Issue 1.
<https://doi.org/10.21428/e3990ae6.483f18da>
- Bygrave, L. A. (2015). *Internet Governance by Contract*. Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199687343.001.0001>
- Carr, C. T., & Hayes, R. A. (2015). Social Media: Defining, Developing, and Divining. *Atlantic Journal of Communication*, 23(1), 46–65. <https://doi.org/10.1080/15456870.2015.972282>
- Citron, D. K., & Wittes, B. (2017). *The Internet Will Not Break: Denying Bad Samaritans Section 230 Immunity* (SSRN Scholarly Paper No. 3007720). <https://papers.ssrn.com/abstract=3007720>
- de Keulenaar, E., Magalhães, J. C., & Ganesh, B. (2023). Modulating moderation: A history of objectionability in Twitter moderation practices. *Journal of Communication*, 73(3), 273–287.
<https://doi.org/10.1093/joc/jqad015>
- Facebook. (2004, June 11). *Terms of Use*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Terms%20of%20Service/20040611_Facebook_TermsOfService.pdf
- Facebook. (2005a, June 30). *Terms of Use*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Terms%20of%20Service/20050630_Facebook_TermsOfService.pdf
- Facebook. (2005b, August 9). *Privacy Policy*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Privacy%20Policy/20050809_Facebook_PrivacyPolicy.pdf
- Facebook. (2007, August 30). *Content Code of Conduct*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Community%20Guidelines/20070830_Facebook_CommunityGuidelines.pdf
- Facebook. (2009, June 3). *Statement of Rights and Responsibilities*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Terms%20of%20Service/20090603_Facebook_TermsOfService.pdf
- Facebook. (2011a, January 27). *Facebook Community Standards*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Community%20Guidelines/20110127_Facebook_CommunityGuidelines.pdf
- Facebook. (2011b, September 15). *Data Use Policy*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Privacy%20Policy/20110915_Facebook_PrivacyPolicy.pdf
- Facebook. (2015, March 18). *Community Standards*. Retrieved from
<https://web.archive.org/web/20150318123123/https://www.facebook.com/communitystandards>
- Facebook. (2018, April 24). *Community Standards*. Retrieved from
<https://web.archive.org/web/20180424131632/https://www.facebook.com/communitystandards/>
- Facebook. (2019, August 3). *Terms of Service*. Platform Governance Archive.
https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Facebook/Terms%20of%20Service/20190803_Facebook_TermsOfService.pdf
- Gillespie, T. (2010). The politics of 'platforms'. *New Media & Society*, 12(3), 347–364.
<https://doi.org/10.1177/1461444809342738>
- Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden*

- decisions that shape social media*. Yale University Press.
- Gollatz, K. (2019). *Die private Governance der Plattformen zur Regelung grenzüberschreitender Kommunikation. Institutionelle Merkmale und die Herausbildung von Regeln im Diskurs*. Universität Zürich.
- Google. (2012, July 27). *Privacy Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Privacy%20Policy/20120727_YouTube_PrivacyPolicy.pdf
- Google. (2018, June 1). *Google Privacy Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Privacy%20Policy/20180601_YouTube_PrivacyPolicy.pdf
- Google. (2023a, June 4). *Updates: Privacy Policy*. <https://policies.google.com/privacy/archive?hl=en-US>
- Google. (2023b, June 4). *Updates: Terms of Service*. <https://policies.google.com/terms/archive?hl=en-US>
- Harvard Kennedy School. (2023). *Fbarchive.org*. <https://fbarchive.org/>
- Helmond, A., & Vlist, F. van der. (2019). Social Media and Platform Historiography: Challenges and Opportunities. *TMG Journal for Media History*, 22(1), Article 1. <https://doi.org/10.18146/tmg.434>
- Horwitz, J. (2021, September 13). *Facebook Says Its Rules Apply to All. Company Documents Reveal a Secret Elite That's Exempt*. WSJ. <https://www.wsj.com/articles/facebook-files-xcheck-zuckerberg-elite-rules-11631541353>
- Instagram. (2011a, March 11). *Instagram*. Retrieved from <https://web.archive.org/web/20110311073313/http://instagram.com/>
- Instagram. (2011b, April 18). *Basic Terms*. Retrieved from <https://web.archive.org/web/20110418202908/http://instagram.com/legal/terms>
- Instagram. (2011c, July 18). *Privacy Policy*. Retrieved from <https://web.archive.org/web/20110718102340/http://instagram.com/legal/privacy>
- Instagram. (2012a, May 3). *Community Guidelines*. Retrieved from <https://web.archive.org/web/20120503130942/http://help.instagram.com/customer/portal/articles/262387-community-guidelines>
- Instagram. (2012b, May 3). *Community Guidelines*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Instagram/Community%20Guidelines/20120503_Instagram_CommunityGuidelines.pdf
- Instagram. (2012c, September 7). *Privacy Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Instagram/Privacy%20Policy/20120907_Instagram_PrivacyPolicy.pdf
- Instagram. (2013, February 3). *Terms of Use*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Instagram/Terms%20of%20Service/20130203_Instagram_TermsOfService.pdf
- Instagram. (2018a, May 15). *Terms of Use*. Retrieved from <https://web.archive.org/web/20180515131230/https://help.instagram.com/581066165581870>
- Instagram. (2018b, July 31). *Data Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Instagram/Privacy%20Policy/20180731_Instagram_PrivacyPolicy.pdf
- Katzenbach, C. (2021). "AI will fix this" – The Technical, Discursive, and Political Turn to AI in Governing Communication. *Big Data & Society*, 8(2). <https://doi.org/10.1177/20539517211046182>
- Kaye, D. (2019). *Speech police: The global struggle to govern the internet*. Columbia Global Reports.
- Kuczerawy, A., & Ausloos, J. (2015). *NoC Online Intermediaries Case Studies Series: European Union and Google Spain*. <https://doi.org/10.2139/ssrn.2567183>
- Meta. (2023a). *Facebook Community Standards*. Transparency Center. <https://transparency.fb.com/policies/community-standards/>

- Meta. (2023b, June 2). *Transparency Center*. <https://transparency.fb.com/>
- Ministry of Europe and Foreign Affairs. (2023). *Open Terms Archive*. <https://opentermsarchive.org/>
- Rajendra-Nicolucci, E. C., & Zuckerman, E. (2021). *An Illustrated Field Guide to Social Media (p. 119)*. Knight First Amendment Institute.
- Rothman, L. (2015, February 4). Happy Birthday, Facebook. *Time*. <https://time.com/3686124/happy-birthday-facebook/>
- Scharlach, R., Hallinan, B., & Shifman, L. (2023). Governing principles: Articulating values in social media platform policies. *New Media & Society*, 14614448231156580. <https://doi.org/10.1177/14614448231156580>
- Siapera, E., & Viejo-Otero, P. (2021). Governing Hate: Facebook and Digital Racism: *Television & New Media*. <https://doi.org/10.1177/1527476420982232>
- Siegler, M. G. (2010, October 6). Instagram Launches With The Hope Of Igniting Communication Through Images. *TechCrunch*. <https://techcrunch.com/2010/10/06/instagram-launch/>
- Suzor, N. P. (2019, June). *Lawless: The Secret Rules That Govern our Digital Lives*. Cambridge Core; Cambridge University Press. <https://doi.org/10.1017/9781108666428>
- The Internet Archive. (2013, September 24). *Wayback Machine APIs*. https://archive.org/help/wayback_api.php
- The Internet Archive. (2023a). *Internet Archive*. <https://archive.org/>
- The Internet Archive. (2023b). *Internet Archive: Wayback Machine*. <https://archive.org/web/>
- The Internet Archive. (2023c). Wayback Machine General Information. *Internet Archive Help Center*. <https://help.archive.org/help/wayback-machine-general-information/>
- The Regents of the University of Michigan. (2023). *Social Media Archive at ICPSR*. Social Media Archive at ICPSR - SOMAR. <https://socialmediaarchive.org>
- Thefacebook. (2004a, February 12). *Welcome to Thefacebook!* Retrieved from <https://web.archive.org/web/20040212031928/http://www.thefacebook.com/>
- Thefacebook. (2004b, June 11). *Terms of Use*. Retrieved from <https://web.archive.org/web/20040611104410/http://www.thefacebook.com/terms.php>
- Thefacebook. (2004c, June 13). *Privacy Policy*. Retrieved from <https://web.archive.org/web/20040613184140/http://www.thefacebook.com/policy.php>
- Twitter. (2006a, September 30). *Twitter*. Retrieved from <https://web.archive.org/web/20060930214639/http://twitter.com/>
- Twitter. (2006b, November 9). *Terms of Service*. Retrieved from <https://web.archive.org/web/20061109073701/http://twitter.com/tos>
- Twitter. (2008a, November 19). *Twitter Privacy Policy*. Retrieved from <https://web.archive.org/web/20081119194405/https://twitter.com/privacy>
- Twitter. (2008b, November 19). *Twitter Privacy Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Privacy%20Policy/20081119_Twitter_PrivacyPolicy.pdf
- Twitter. (2009a, January 18). *The Twitter Rules*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Community%20Guidelines/20090118_Twitter_CommunityGuidelines.pdf
- Twitter. (2009b, January 18). *The Twitter Rules*. Retrieved from <https://web.archive.org/web/20090118211301/http://twitter.zendesk.com/forums/26257/entries/18311>
- Twitter. (2009c, September 17). *Previous Terms of Service*. Retrieved from <https://web.archive.org/web/20090917045821/https://twitter.com/tos/previous>
- Twitter. (2009d, September 18). *Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Terms%20of%20Service/20090918_Twitter_TermsOfService.pdf
- Twitter. (2010, March 1). *The Twitter Rules*. Retrieved from <https://web.archive.org/web/20100301141609/http://twitter.zendesk.com/forums/26257/entries/18311>

- Twitter. (2013, May 30). *The Twitter Rules*. Retrieved from <https://web.archive.org/web/20130530082742/https://support.twitter.com/articles/18311>
- Twitter. (2016, October 5). *Twitter Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Terms%20of%20Service/20161005_Twitter_TermsOfService.pdf
- Twitter. (2017a, June 16). *Previous Privacy Policies*. Retrieved from <https://web.archive.org/web/20170616021110/https://twitter.com/en/privacy/previous>
- Twitter. (2017b, December 6). *Rules and policies*. Retrieved from <https://web.archive.org/web/20171206081009/https://help.twitter.com/en/rules-and-policies>
- Twitter. (2019, June 7). *The Twitter Rules*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Community%20Guidelines/20190607_Twitter_CommunityGuidelines.pdf
- Twitter. (2020a, January 6). *Twitter Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Terms%20of%20Service/20200106_Twitter_TermsOfService.pdf
- Twitter. (2020b, June 19). *Twitter Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/Twitter/Terms%20of%20Service/20200619_Twitter_TermsOfService.pdf
- Twitter. (2023a, June 4). *Previous Terms of Service*. <https://twitter.com/content/twitter-com/legal/en/tos/previous>
- Twitter. (2023b, June 5). *The Twitter Rules*. <https://help.twitter.com/en/rules-and-policies/twitter-rules>
- YouTube. (2005a, April 28). *Privacy Policy*. Retrieved from <https://web.archive.org/web/20050428195502/http://www.youtube.com/privacy.php>
- YouTube. (2005b, April 28). *Privacy Policy*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Privacy%20Policy/20050428_YouTube_PrivacyPolicy.pdf
- YouTube. (2005c, April 28). *Terms of Use*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Terms%20of%20Service/20050428_YouTube_TermsOfService.pdf
- YouTube. (2005d, May 5). *YouTube*. Retrieved from <https://web.archive.org/web/20050505071842/http://www.youtube.com/>
- YouTube. (2008a, November 12). *YouTube Community Guidelines*. Retrieved from https://web.archive.org/web/20081112004550/www.youtube.com/t/community_guidelines
- YouTube. (2008b, November 20). *YouTube Privacy Notice*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Privacy%20Policy/20090320_YouTube_PrivacyPolicy.pdf
- YouTube. (2009, March 20). *Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Terms%20of%20Service/20090320_YouTube_TermsOfService.pdf
- YouTube. (2011, May 25). *Thanks, YouTube community, for two BIG gifts on our sixth birthday!* YouTube Official Blog. <https://blog.youtube/news-and-events/thanks-youtube-community-for-two-big/>
- YouTube. (2018, June 2). *Terms of Service*. Platform Governance Archive. https://github.com/PlatformGovernanceArchive/pga-corpus/blob/main/Versions/PDF/YouTube/Terms%20of%20Service/20180602_YouTube_TermsOfService.pdf
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (1st ed.). PublicAffairs.