Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation

Victor Le Pochat
imec-DistriNet, KU Leuven
Belgium
victor.lepochat@kuleuven.be

Tom Van Goethem
imec-DistriNet, KU Leuven
Belgium
tom.vangoethem@kuleuven.be

Samaneh Tajalizadehkhoob
ICANN
United States
samaneh.tajali@icann.org

Maciej Korczyński
Univ. Grenoble Alpes, CNRS,
Grenoble INP, LIG
France
maciej.korczynski@univ-grenoble-alpes.fr

Wouter Joosen
imec-DistriNet, KU Leuven
Belgium
wouter.joosen@kuleuven.be

ABSTRACT

The web security and privacy community often relies on top sites rankings to measure prevalence across a representative sample of websites, contextualize findings by website popularity, or classify benign and malicious domains. However, prior work showed that previously available rankings had undesirable properties for research. As a service to the community, for nearly four years, we have provided the *Tranco* ranking, a ranking with a strong focus on improving upon properties important for research, such as reproducibility and transparency. In this short paper, we describe how Tranco combines existing rankings to generate a more reliable ranking, and show how over 350 research studies as well as several high-profile entities in industry, the media, and beyond have all used Tranco as the basis for their web measurements.

1 INTRODUCTION

Rankings of the most popular or "top" domain names have long been a prominent data source in web security and privacy as well as Internet measurement research, where they serve multiple purposes. First, they allow to select a meaningful sample of websites for measuring the prevalence of a certain phenomenon or vulnerability, or evaluate a novel technique, attack, or defense. With an estimated 350 million registered domain names [44], it is often infeasible to study all websites on the Internet, in terms of the resources required, in particular when collecting data by visiting web pages through an instrumented browser. For example, in one recent study, crawling 50 pages on 100,000 websites using Chromium took 8 days across 60 parallel crawlers [30]. Many of these websites and domain names are also relatively uninteresting: they may be special-purpose domains that are unlikely to be visited by the public at large. Selecting only relatively popular sites provides a more representative view on what end users may encounter while they browse the web. Second, the relative ranking of domains also allows to contextualize research findings, indicating whether the popularity of a website correlates with the occurrence of a certain phenomenon. Third, popular websites are sometimes considered as a source of benign domains, e.g., in a classifier of benign and malicious domains, although this is not a necessary or given property of these rankings.

Finalist submission for the 2022 ACSAC Cybersecurity Artifacts Impact Award.

These rankings form an essential part of many web and Internet research studies: in our 2018 literature survey [23], we found 133 papers in top-tier security conferences across four years to use at least one ranking, and over 350 studies have used our Tranco ranking so far.

Historically, major commercial players published top sites rankings. In the SEO and marketing space, these were Alexa (since 2008), Quantcast (since 2007), and later Majestic (since 2012). (Cisco) Umbrella also published a ranking (since 2016), as a large DNS resolver. Since then, the space of commercial top sites rankings has further evolved. The Quantcast ranking was silently discontinued in April 2020. The most commonly used Alexa ranking is also slated to disappear in December 2022 [40].

Analyzing domain rankings. Despite the importance of these commercial rankings to research, and decades of use, the research community only scrutinized them and became fully aware of the issues surrounding them from 2018 onward. Two contemporaneous works studied these rankings in depth: the *"A long way to the top"* paper by Scheitle et al. [36], and our Tranco paper [23]. They reverse engineered the proprietary and opaque methods used to construct the four commercial rankings, surveyed their usage in research usage, measured characteristics such as stability and similarity, and quantified the potential impact on Internet measurement and security research respectively. In our Tranco work, we also showed that all four rankings could be manipulated to insert any domain.

2 TRANCO: IMPROVING DOMAIN RANKINGS

In 2019, we first released our Tranco list, a new domain ranking for use in Internet and web measurements. We designed our ranking with a strong focus on improving upon properties important for research, such as reproducibility and transparency. In our original paper [23], we transparently describe the construction method for the Tranco list. We also replicate this description on the Tranco website.[1] The code for constructing lists is also openly sourced[2] to provide full transparency and to allow for building upon our work.

[1] https://tranco-list.eu/methodology
[2] https://github.com/DistriNet/tranco-list

List construction. In general, lists are constructed by averaging the ranks of domains across a chosen period of time and a set of external ranking providers, with the possibility to apply filters. This method has the goal of enhancing the research-oriented properties of our list, such as stability and similarity. Moreover, we provide adaptability for researchers, by giving them the ability to customize the list to their needs, such as selecting the most appropriate providers (e.g., only those who primarily rank websites), selecting an appropriate period of time (e.g., a longer period for more stability), or applying appropriate filters (e.g., to remove known malicious websites).

Ranks are averaged over the lists of one or more providers, as available at the time of construction. The Umbrella and Majestic have always been available, the Alexa list is currently available but slated to disappear [40], the Quantcast list has already been retired, and the Farsight list was added in 2022 but only for inclusion in the default list. We provide two combination methods for averaging: for a list of length N, the Borda count scores items with $N, N-1, ..., 1, 0$ points; the Dowdall rule scores items with $1, 1/2, ..., 1/(N-1), 1/N$ points [11]. The latter reflects the observation of Zipf's law distribution in website traffic [2, 9]. Once this averaging method has been selected, the scoring is applied to each source list individually, after which these scores are summed and the domains are ranked on this aggregate score, yielding the final ranking. When averaging over D days for the lists of P providers, this means $D \times P$ lists are used to compute the final Tranco score and rank.

We add the possibility to apply filters for creating a list containing domains that align more closely with the intended purpose of the list. The first set of filters heuristically targets sites that are more likely to be (dis)similar. A researcher can select only domains with certain TLDs, to favor sites that are more likely to be associated with particular countries or sectors (e.g., government domains), or exclude (overly represented) TLDs. A researcher can also limit the set of pay-level domains that have the same 'second-level' label and only differ in TLD (e.g., `google.com`, `google.de`, `google.fr`, ...) to only the (one) top-ranked domain, to reduce the dominance of particular organizations. Finally, a researcher can select only certain subdomains, e.g. searching for `login.*` subdomains to obtain a list of likely authentication services.

The second set of filters increases the likelihood that the domains in the list have certain properties. By default, we merge domains from both web-oriented and DNS-based lists. Because of the latter, the resulting list also contains infrastructural domains such as CDNs, nameservers. While these can certainly be viewed as popular domains, they are less useful for studies of websites, as they are unlikely to serve any meaningful web content, in particular on their root page. A researcher can select actual websites that additionally are regularly visited through a major browser by filtering on domains included in Google's Chrome User Experience Report, a set of over 16 million distinct domains that are said to be 'popular destinations on the web' [1] and have been found to reflect popularity well [32]. Next to that, if a researcher wants to assume that ranked popular domains are benign, they can remove known malicious domains flagged by Google Safe Browsing [33]. Researchers can also filter out domains that appear only on one or a few lists and/or days, to avoid domains that are only marked as popular by one provider: these may point to isolated manipulation.

To make using Tranco as easy as possible, reduce the choices that a researcher has to make, and encourage standardization, we set default options for all customization parameters and use them when generating the daily main Tranco ranking. Our standard list applies the Dowdall scoring rule, chosen to incorporate the Zipf's law distribution, to all lists available at the given time, which currently are Alexa, Farsight, Majestic, and Umbrella. To improve stability by smoothing out changes over time, we default to using the lists from the past 30 days. We do not apply any other filters such that the default list can cover as diverse a set of use cases as possible. Our long-term evaluation showed that these default parameters result in a stable, robust and comprehensive ranking [22].

List and data availability. While the list generation method described in our paper can be applied by researchers themselves on any set of rankings that they have at their disposal, we also make an online service available at https://tranco-list.eu, where researchers can generate customized lists using our archive of external rankings, as well as download the standard daily Tranco ranking with the parameters described in the previous paragraph. This service has been online since the publication of the Tranco paper in February 2019, i.e., for almost 4 years, providing a sustained long-term service to our research community.

On this website, we provide multiple ways to access the Tranco data, with ease of use in mind. The main list and any customized list is formatted as a CSV that serves as a drop-in replacement for the Alexa and Umbrella lists –which were previously commonly used– to minimize the development effort for researchers to change to the Tranco list in both ongoing and new experiments. Moreover, the most recent daily list is available from a stable and well-known URL, facilitating automatically retrieving the most recently updated list on a regular basis. For bulk queries across the entire history of daily Tranco lists, we also upload the daily list to Google Cloud's BigQuery, where the historical data set can be queried using SQL.

We also provide an API where users can retrieve list metadata and historical ranks for domains. The latter allows to quickly track the evolution of a domain's popularity without needing to download all previous lists and filter on that domain. Authenticated users can also use the API to generate customized lists programmatically. We have also developed a Python package that users can integrate with their Python programs to download, query, and generate lists.

Reproducibility. Studies rarely mention the date on which a ranking was retrieved [10, 23, 36], and even if they would, it is hard to obtain any other ranking than the most recent one. The reproducibility of these studies is therefore harmed, as the exact set of examined websites cannot be retrieved afterwards.

From the beginning, Tranco was developed with the goal to support and improve reproducibility. Core to this is the permanent archival of any daily or customized list. Every list receives a permanent ID, and through a link that contains this ID, anyone can retrieve the exact set of domains in that list. The page behind this link also contains a permanent record of the exact configuration settings of a list, such as the providers used and the period over which lists were averaged. To encourage researchers to give a precise description of the data sets they used in their study, this page also contains a suggested citation to refer to the generated list and the Tranco project. By making the historical lists more easily

Table 1: Distribution of papers citing Tranco across venue types and years, across the 250 papers for which we were able to retrieve bibliographic data. * Until August 2022.

	2019	2020	2021	2022*	Total
Top-tier security conferences	1	7	15	12	35
co-located workshops	6	3	6	4	19
Other security conferences	6	9	12	11	38
Privacy conferences	0	0	9	3	12
Web & Internet measurement	9	14	13	10	46
Other conferences	6	9	12	8	35
Journals	2	6	9	11	28
Unknown	3	7	11	16	37
Total	33	55	87	75	250

accessible, we improve the ability to replicate any studies that use them, contributing to sound research practices.

3 IMPACT

Academic impact. Tranco has demonstrated a large impact on web and Internet measurement research in both academia and beyond. The most readily measurable impact is the number of citations in academic research papers, currently amounting to over 350 in less than 4 years, according to Google Scholar[3], therefore being one of the top 10 cited security papers from 2019[4]. Tranco has been used in publications at the four top-tier security conferences and their co-located workshops, other security and privacy conferences, as well as the prominent Web and Internet measurement conferences (Table 1). Tranco has also found its way into other research domains, e.g., being used in the process of creating neural models for information extraction from websites [34].

Tranco has arguably become the *de facto* standard top list in the research field. Demir et al. [10] highlighted Tranco as a "work[] that aim[s] to provide best practices" and that "ha[s] a positive impact on our community." Anecdotally, we know reviewers are aware of this position. We have received comments from reviewers either suggesting we use Tranco instead of Alexa or commending us for using Tranco, as well as pointing out the reproducibility harm when omitting the exact Tranco list used in a study. Generally, we see that papers mention the Tranco project and the exact list ID well, providing an immediate benefit to their reproducibility and validity.

Industry and media impact. The Tranco list is used by prominent industry players, showing its wide applicability. Tranco is a contributing data partner to projects by the Electronic Frontier Foundation (for their 'Privacy Badger' extension) [27], the Internet Society's Pulse project (measuring the health, availability and evolution of the Internet) [20], the Brave browser (filtering search results and minimizing filter lists) [7, 37], ScamAdviser (as a popularity indicator),[5] URLhaus by abuse.ch (sharing malware URLs) [42], BuiltWith [8], and W3Techs [13] (the latter two measuring web technology usage). Tranco has been used for measurements by

[3]https://scholar.google.be/scholar?oi=bibs&cites=14996934840507597 6,10234769677796230547,17897712023882147302

[4]https://www.sec.tu-bs.de/~konrieck/topnotch/sec_2010s.html

[5]see, e.g., https://www.scamadviser.com/check-website/scamadviser.com

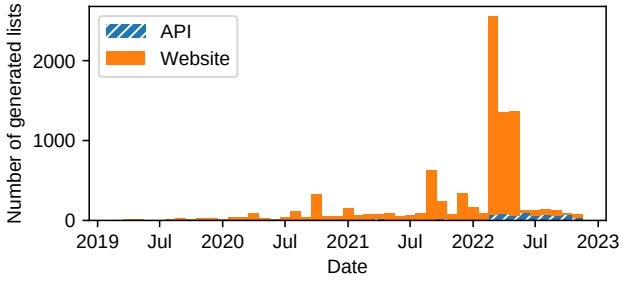

Figure 1: Number of customized lists generated over time, through the website or the API.

Mozilla [38], Cloudflare [45], F5 Labs [47], Palo Alto Networks' Unit 42 [24], Avast [6], and ICANN [17, 18]. Tranco is available as a filter list in the threat intelligence platforms from MISP [28], SEKOIA [41], Intel Owl [5], and ThreatConnect [14].

On the media side, Tranco was used by journalists from the Norwegian public broadcaster NRK [15] and The Markup [26, 35], and was mentioned in news articles from The Verge [46] and SecurityWeek [21] that reported on the demise of the Alexa ranking. Prominent security researchers Scott Helme and Troy Hunt use Tranco for their scans of the state of security on the web [16] and the Why No HTTPS? project [19], respectively.

Usage statistics. The online service for accessing the daily standard Tranco ranking and generating customized lists is publicly and freely available. Since the Tranco ranking was made public in February 2019 until 21 November 2022, 1,374 daily lists were published on this service. In addition, 497 distinct users generated 9,268 customized lists, with a noticeable uptick for the latter in 2022 (Figure 1). The API is also increasingly popular for generating lists, having been used for 712 lists already since early 2022. Filters are used in moderation, with 858 lists selecting probable websites using the Chrome User Experience Report, 399 lists including or excluding certain TLDs, 346 lists removing known malicious websites using Google Safe Browsing, and 262 lists filtering organizations.

According to a one-week log from September 2022 of the AWS S3 bucket serving the daily list, around 25 unique endpoints request the daily list per hour, with a noticeable spike to around 70 endpoints when the new daily list is released at midnight (Figure 2). Based on web analytics data available to us, the Tranco website receives at least 700 daily unique visitors, and over 1,000 monthly clickthroughs from Google search results. These last metrics ignore those who opt out of analytics tracking, as well as those who directly access list ZIP files, in particular the daily list, which is available at a stable URL that does not require passing through the Tranco website.

The GitHub repository with our open-source implementation has been forked 12 times and starred 74 times, although this underrepresents the impact, as not many people will need to consult the list generation code but instead use the end product that is available on our website. More telling is the usage of the Tranco Python package, which has been downloaded over 121,000 times [3], with a noticeable increase in fall 2021 (Figure 3). This shows that our open dataset is integrated into a variety of software projects.

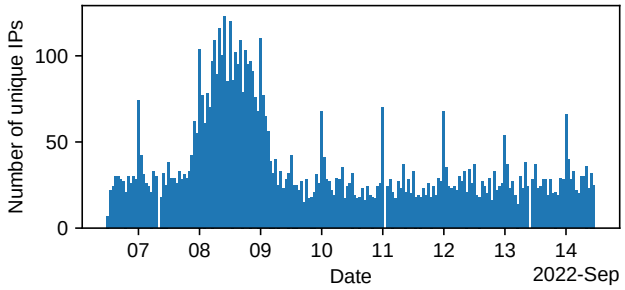

Figure 2: Number of unique IPs downloading the daily list, grouped by hour.

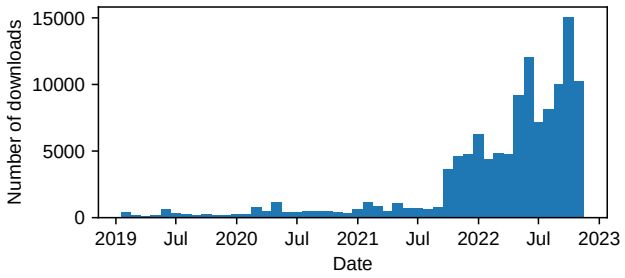

Figure 3: Number of monthly downloads of the Tranco Python package.

Follow-up work. As a final testament to the impact of our research and the artifact that resulted from it, and the importance of having such a well validated dataset available, we discuss other research and industry projects that have proposed and published rankings similar to Tranco.

Naab et al. [31] developed and published prefix top lists[6] that derive a ranking of network prefixes from existing domain rankings, intended for measurements of Internet infrastructure such as CDNs or core routers. Aqeel et al. [4] developed and published Hispar,[7] a ranking that includes internal pages, i.e., pages beyond the landing pages that are typically measured when visiting the websites in a domain ranking. They discover these internal pages through search engine results. The ranking is intended for increasing the per-website coverage of a measurement, and improve the validity of findings for a website. Unfortunately, both initiatives appear to have stopped publishing new rankings.

In 2022, Xie et al. [48] developed the SecRank ranking method for passive DNS traffic, which is based on weighted voting across the domain preferences of individual IP addresses. They publish a list where they apply their method to passive DNS data from 114 DNS (a Chinese public DNS resolver),[8] but downloading this list requires an account and is therefore more difficult to implement automatically.

[6]https://prefixtoplists.net.in.tum.de/
[7]https://hispar.cs.duke.edu/
[8]https://secrank.cn/topdomain

In industry, there are two prominent new initiatives. DomainTools published a passive DNS-based ranking using data from their Farsight network [12]. Within their threat intelligence product Iris, they also incorporate a ranking that is aggregated using the same method as Tranco [12]. Cloudflare also published their 'Radar' passive DNS-based ranking [25]. Other companies have also published independent analyses of Tranco. Infoblox compared their (proprietary and not publicly available) InfoRanks ranking [39] to existing rankings, including Tranco, on properties such as similarity and benignness. DeepSee analyzed the influence of individual component rankings, and in particular Alexa, on the final ranks in Tranco [29].

4 OUTLOOK TO THE FUTURE

We currently already provide a robust, well used, and increasingly important service to the community. However, we still have many exciting plans for new tools that would make Tranco even more useful, and would guarantee its long-term availability and impact.

We continue monitoring the space of top sites lists, and track both the arrival of new lists and the discontinuation of existing lists. Concretely, we are looking into integrating the Chrome User Experience Report [1] and Cloudflare Radar [25] rankings, after necessary recalibration due to their bucketed ranks. We have ongoing work designing a domain ranking using passive DNS data in a privacy-preserving manner. This ranking would improve upon existing rankings in terms of transparency, as the ranking method would be fully transparent; and availability, as it would depend on a raw data source that is unlikely to disappear soon.

The current Tranco list incorporates rankings across a variety of vantage points, sourced from both web-based measurements and DNS traffic. Since studies usually consider either only websites or all types of Internet infrastructure, having the combination of all vantage points in one ranking may not be desirable, especially for the web-focused studies. While we already provide the ability to select only web-focused rankings for generating a customized list, we are developing a service classification that would indicate whether a domain rather hosts a website or is used as part of broader Internet infrastructure. We then also intend to generate separate rankings for web and infrastructural resources, providing researchers the ability to select the ranking that is most appropriate for their research. We are also exploring solutions in the space of domain classification based on openly available data sets, as researchers currently rely on commercial solutions that are difficult to access at scale and may be insufficient in accuracy and coverage [43].

For the Tranco website, we are working to develop a dashboard through which researchers can visually explore the Tranco ranking to better understand its properties, such as the long-term stability, relative importance of the component rankings, or composition in terms of, e.g., TLDs or organizations. Finally, to increase the resilience of Tranco, ensure that the previously generated lists remain retrievable and therefore ensure that the research that used those lists remains maximally reproducible, we intend to duplicate the lists on a long-term stable (research) repository such as OSF[9] or the Internet Archive. With these features gradually being added to Tranco, we guarantee that it remains available, useful, relevant, and growing in impact for the foreseeable future.

[9]https://osf.io/

ACKNOWLEDGMENTS

This research is partially funded by the Research Fund KU Leuven, and by the Flemish Research Programme Cybersecurity. Victor Le Pochat was supported by a PhD Fellowship of the Research Foundation Flanders - FWO (11A3419N, 11A3421N).

REFERENCES

[1] 2022. About CrUX. Google Inc. (June 23, 2022). https://developer.chrome.com/docs/crux/about/.

[2] Lada A. Adamic and Bernardo A. Huberman. 2002. Zipf's law and the Internet. *Glottometrics*, 3, 143–150.

[3] 2022. Analyzing PyPI package downloads. Python Packaging User Guide. https://packaging.python.org/en/latest/guides/analyzing-pypi-package-downloads/.

[4] Waqar Aqeel, Balakrishnan Chandrasekaran, Anja Feldmann, and Bruce M. Maggs. 2020. On Landing and Internal Web Pages: The Strange Case of Jekyll and Hyde in Web Performance Measurement. In *2020 ACM Internet Measurement Conference* (IMC '20), 680–695. DOI: 10.1145/3419394.3423626.

[5] Eshaan Bansal. 2020. Intel Owl Release v1.0.0. The Honeynet Project. (July 5, 2020). https://www.honeynet.org/2020/07/05/intel-owl-release-v1-0-0/.

[6] Joe Bosso. 2022. How does your location affect your online privacy? Avast. (Jan. 5, 2022). https://blog.avast.com/location-and-online-privacy-avast.

[7] 2022. Brave Search passes 2.5 billion queries in its first year, and debuts Goggles feature that allows users to choose their own search rankings. Brave. (June 22, 2022). https://brave.com/search-anniversary/.

[8] Gary Brewer. 2018. What is Tranco, Quantcast, Majestic and Umbrella Numbers? BuiltWith. (Aug. 10, 2018). https://kb.builtwith.com/general-questions/what-is-tranco-quantcast-majestic-and-umbrella-numbers/.

[9] Aaron Clauset, Cosma Rohilla Shalizi, and Mark E. J. Newman. 2009. Power-Law Distributions in Empirical Data. *SIAM Review*, 51, 4, 661–703. DOI: 10.1137/070710111.

[10] Nurullah Demir, Matteo Große-Kampmann, Tobias Urban, Christian Wressnegger, Thorsten Holz, and Norbert Pohlmann. 2022. Reproducibility and Replicability of Web Measurement Studies. In *ACM Web Conference 2022* (WWW '22), 533–544. DOI: 10.1145/3485447.3512214.

[11] Jon Fraenkel and Bernard Grofman. 2014. The Borda Count and its real-world alternatives: Comparing scoring rules in Nauru and Slovenia. *Australian Journal of Political Science*, 49, 2, 186–205. DOI: 10.1080/10361146.2014.900530.

[12] Aaron Gee-Clough. 2022. Mirror, Mirror, on the Wall, Who's the Fairest (website) of Them all? DomainTools. (Apr. 28, 2022). https://www.domaintools.com/resources/blog/mirror-mirror-on-the-wall-whos-the-fairest-website-of-them-all.

[13] Matthias Gelbmann. 2020. Extension of our website sample set. W3Techs. (Mar. 31, 2020). https://w3techs.com/blog/entry/extension_of_our_website_sample_set.

[14] Drew Gidwani. 2020. Herding Cattle: ThreatConnect's Vision for Better Intel Feeds. ThreatConnect. (Feb. 19, 2020). https://threatconnect.com/blog/cal-2-4-introducing-cal-feeds/.

[15] Martin Gundersen. 2019. Dette er de norske nettsidene som sporer deg mest. Norwegian. NRKbeta. (May 28, 2019). https://nrkbeta.no/2019/05/28/dette-er-de-norske-nettsidene-som-sporer-deg-mest/.

[16] Scott Helme. 2019. Top 1 Million Analysis - September 2019. (Oct. 3, 2019). https://scotthelme.co.uk/top-1-million-analysis-september-2019/.

[17] Paul Hoffman. 2021. Collecting "Typical" Domain Names for Web Servers. Tech. rep. OCTO-023. (Feb. 24, 2021). https://www.icann.org/en/system/files/files/octo-023-24feb21-en.pdf.

[18] Paul Hoffman. 2022. DNSSEC Algorithm Use in 2022. Tech. rep. OCTO-033. (Apr. 4, 2022). https://www.icann.org/en/system/files/files/octo-033-04apr22-en.pdf.

[19] Troy Hunt. 2021. Why No HTTPS? (Aug. 12, 2021). https://whynohttps.com/.

[20] 2021. Internet Society Pulse – Partners. Internet Society. https://pulse.internetsociety.org/partners.

[21] Eduard Kovacs. 2022. Amazon's Shuttering of Alexa Ranking Service Hits Cybersecurity Industry. *SecurityWeek*, (May 6, 2022). https://www.securityweek.com/impact-alexa-ranking-service-shutdown-cybersecurity-industry.

[22] Victor Le Pochat, Tom Van Goethem, and Wouter Joosen. 2019. Evaluating the Long-term Effects of Parameters on the Characteristics of the Tranco Top Sites Ranking. In *12th USENIX Workshop on Cyber Security Experimentation and Test* (CSET '19). https://www.usenix.org/conference/cset19/presentation/lepochat.

[23] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *26th Annual Network and Distributed System Security Symposium* (NDSS '19). DOI: 10.14722/ndss.2019.23386.

[24] Daiping Liu and Ruian Duan. 2021. Dangling Domains: Security Threats, Detection and Prevalence. Palo Alto Networks – Unit 42. (Sept. 16, 2021). https://unit42.paloaltonetworks.com/dangling-domains/.

[25] Celso Martinho and Sabina Zejnilovic. 2022. Goodbye, Alexa. Hello, Cloudflare Radar Domain Rankings. The Cloudflare Blog. (Sept. 30, 2022). https://blog.cloudflare.com/radar-domain-rankings/.

[26] Surya Mattu and Aaron Sankin. 2020. How We Built a Real-time Privacy Inspector. *The Markup*, (Sept. 22, 2020). https://themarkup.org/blacklight/2020/09/22/how-we-built-a-real-time-privacy-inspector.

[27] Alexei Miagkov, Andrés Arrieta, and Bennett Cyphers. 2018. Giving Privacy Badger a Jump Start. Electronic Frontier Foundation. (Aug. 22, 2018). https://www.eff.org/deeplinks/2018/08/giving-privacy-badger-jump-start.

[28] [SW], misp-warninglists 2022. MISP. URL: https://github.com/MISP/misp-warninglists.

[29] Rocky Moss and Edward Krueger. 2022. We Bid Goodbye to Alexa Rankings, and Measure Its Contribution to the Tranco List (Pre-May). DeepSee. (May 25, 2022). https://deepsee.io/blog/we-bid-goodbye-to-alexa-rankings-and-measure-its-contribution-to-the-tranco-list-pre-may.

[30] Marius Musch, Robin Kirchner, Max Boll, and Martin Johns. 2022. Server-Side Browsers: Exploring the Web's Hidden Attack Surface. In *2022 ACM Asia Conference on Computer and Communications Security* (ASIA CCS '22), 1168–1181. DOI: 10.1145/3488932.3517414.

[31] Johannes Naab, Patrick Sattler, Jonas Jelten, Oliver Gasser, and Georg Carle. 2019. Prefix Top Lists: Gaining Insights with Prefixes from Domain-based Top Lists on DNS Deployment. In *2019 Internet Measurement Conference* (IMC '19), 351–357. DOI: 10.1145/3355369.3355598.

[32] Kimberly Ruth, Deepak Kumar, Brandon Wang, Luke Valenta, and Zakir Durumeric. 2022. Toppling Top Lists: Evaluating the Accuracy of Popular Website Lists. In *22nd ACM Internet Measurement Conference* (IMC '22), 374–387. DOI: 10.1145/3517745.3561444.

[33] [n. d.] Safe Browsing. Google Inc. https://safebrowsing.google.com/.

[34] Aidan San, Jan Bakus, Colin Lockard, David Ciemiewicz, Yangfeng Ji, Sandeep Atluri, Kevin Small, and Heba Elfardy. 2022. PLAtE: A Large-scale Dataset for List Page Web Extraction. (2022). arXiv: 2205.12386. DOI: 10.48550/ARXIV.2205.12386.

[35] Aaron Sankin and Surya Mattu. 2020. The High Privacy Cost of a "Free" Website. *The Markup*, (Sept. 22, 2020). https://themarkup.org/blacklight/2020/09/22/blacklight-tracking-advertisers-digital-privacy-sensitive-websites.

[36] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D. Strowes, and Narseo Vallina-Rodriguez. 2018. A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists. In *2018 Internet Measurement Conference* (IMC '18), 478–493. DOI: 10.1145/3278532.3278574.

[37] [SW], Slim List System 2022. Brave. URL: https://github.com/brave/slim-list-lambda.

[38] Mike Taylor. 2019. TLS 1.0 and 1.1 Removal Update. Mozilla Hacks. (May 15, 2019). https://hacks.mozilla.org/2019/05/tls-1-0-and-1-1-removal-update/.

[39] Laura Teixeira da Rocha. 2021. InfoRanks: Infoblox Ranking Service. Infoblox. https://www.infoblox.com/wp-content/uploads/infoblox-whitepaper-inforanks-infoblox-ranking-service.pdf.

[40] The Alexa.com Team. 2021. We will be retiring Alexa.com on May 1, 2022. Alexa Internet, Inc. (Dec. 22, 2021). https://support.alexa.com/hc/en-us/articles/4410503838999.

[41] 2022. Tranco. SEKOIA.IO. https://docs.sekoia.io/tip/features/automate/library/tranco/.

[42] 2022. URLhaus API. abuse.ch. https://urlhaus.abuse.ch/api/.

[43] Pelayo Vallina, Victor Le Pochat, Álvaro Feal, Marius Paraschiv, Julien Gamba, Tim Burke, Oliver Hohlfeld, Juan Tapiador, and Narseo Vallina-Rodriguez. 2020. Mis-shapes, Mistakes, Misfits: An Analysis of Domain Classification Services. In *2020 Internet Measurement Conference* (IMC '20), 598–618. DOI: 10.1145/3419394.3423660.

[44] Verisign, (Ed.) 2022. The Domain Name Industry Brief 19 2 (June 2022). https://www.verisign.com/assets/domain-name-report-Q12022.pdf.

[45] Tanya Verma and Sudheesh Singanamalla. 2020. Improving DNS Privacy with Oblivious DoH in 1.1.1.1. Cloudflare. (Dec. 8, 2020). https://blog.cloudflare.com/oblivious-dns/.

[46] James Vincent. 2021. Amazon is retiring Alexa – no, not that one. *The Verge*, (Dec. 9, 2021). https://www.theverge.com/2021/12/9/22825744/amazon-retiring-alexa-web-ranking-sevice.

[47] David Warburton. 2021. The 2021 TLS Telemetry Report. F5 Labs. (Oct. 20, 2021). https://www.f5.com/labs/articles/threat-intelligence/the-2021-tls-telemetry-report.

[48] Qinge Xie, Shujun Tang, Xiaofeng Zheng, Qingran Lin, Baojun Liu, Haixin Duan, and Frank Li. 2022. Building an Open, Robust, and Stable Voting-Based Domain Top List. In *31st USENIX Security Symposium* (USENIX Security '22), 625–642.