Technical Profiling of Blocked Domains in Parental Control Systems

Yousef Gouriye
University of Twente
Enschede, Netherlands

1 ABSTRACT

Parental control tools are used to block domains that may be inappropriate for children. However, these tools often do not clearly explain how they decide which domains to block. This project will investigate whether blocked domains share technical features, such as their top-level domain, hosting location, or other hidden technical details that could suggest hidden rules or biases in how filtering is done. Using a dataset of websites blocked by different parental control systems, the goal is to find patterns in technical attributes that can help us understand what influences these systems' decisions beyond just website content. Our analysis reveals that blocking decisions are not only driven by content categories but also by technical factors like IP subnets, WHOIS metadata, and domain age. We found that recently created domains and domains from large hosting providers are more likely to be blocked. Additionally, entire IP subnets were often blocked regardless of their actual content, highlighting potential infrastructural bias in PCS filtering behavior.

2 INTRODUCTION

In the current digital landscape, parental control systems (PCS) have become increasingly essential for managing children's exposure to harmful or inappropriate online content. These tools, integrated into routers, DNS providers, operating systems, and third-party software, aim to enforce safe browsing environments by blocking access to domains that violate predefined content standards. As internet usage among minors continues to rise—driven by educational tools, social media, and entertainment platforms—there is growing pressure on parents, educators, and institutions to implement measures that filter unsafe content and protect children online. Recent surveys report that 86% of UK parents with children under 11 have used at least one parental control setting, with adoption rates remaining high among parents of older children [6, 11, 12]. Consequently, PCS technologies have evolved from simple keyword-based filters—often limited and prone to overblocking—to complex systems involving real-time URL classification, reputation scoring, and domain-level blocking. Modern solutions increasingly use machine learning and context-aware analysis to improve filtering accuracy [1].

While these systems are marketed as content-based filters, offering protection by evaluating the textual or categorical makeup of websites, there is increasing suspicion and partial evidence that they also rely on technical metadata when making blocking decisions. This includes characteristics like the domain's IP address, the registrar, WHOIS information, hosting provider, or even creation and update dates—factors that are typically unrelated to the actual content being accessed. These heuristics may stem from operational shortcuts, shared infrastructure blacklists, or efforts to preemptively block domains associated with risky behavior. However, the opacity of these systems raises critical questions about fairness, overblocking, collateral damage, and accountability.

The lack of transparency in PCS decision-making is a growing concern. While some filtering outcomes are intuitive (e.g., adult or gambling sites), others appear arbitrary—raising the possibility that seemingly innocuous or legitimate domains are blocked due to shared hosting or registrar choices. These false positives not only frustrate users but may also suppress access to educational, political, or social resources. Moreover, users are often given no explanation or recourse when a site is blocked, and many vendors do not disclose the heuristics behind their filtering logic. Studies by digital rights organizations highlight that overblocking can prevent access to educational or advocacy websites, and that appeals processes are rarely available [10].

A promising but underexplored research direction is the role of technical domain features in influencing blocking behavior. Previous work by researchers such as Liberato et al. [8] provided valuable insights into the classification accuracy and inconsistencies of PCS tools, but primarily focused on content categories. Our study extends this work by investigating non-content-based factors—what we refer to as "technical profiling" of domains. Technical profiling involves the analysis of structural and metadata-level features, such as the domain's IP subnet (/24 blocks), WHOIS field completeness, registrar brand, DNS configurations, and registration timelines, to identify whether certain patterns correlate with an increased likelihood of blocking.

Understanding these patterns is not only academically valuable but also socially and technically important. It can reveal unintentional biases in filtering systems, identify systemic vulnerabilities (e.g., reliance on shared infrastructure), and guide developers and policymakers in creating more transparent and equitable internet safety technologies.

To carry out this investigation, we leverage a robust dataset comprised of blocked domain lists from seven real-world PCS implementations and a baseline sample from the Tranco top-1M domain list. Each domain has been classified using Cisco Umbrella's category taxonomy and enriched with technical metadata, including WHOIS records and DNS data. This allows us to analyze patterns both within and across PCS platforms.

The central research questions guiding this study are as follows:

- To what extent do technical attributes—such as IP addresses, domain registrars, and WHOIS metadata—correlate with the likelihood of a domain being blocked?
- Are there consistent cross-system patterns that suggest shared heuristics or infrastructural biases among PCS tools?
- Do specific technical configurations trigger blocking regardless of a domain's actual content?

TScIT 43, July 4, 2025, Enschede, The Netherlands

To answer these questions, we formulate the following objectives:

(1) Collect and analyze technical metadata from both blocked and unblocked domains to identify correlations with blocking behavior.
(2) Evaluate similarities and differences in blocking patterns across multiple PCS systems, focusing on infrastructural clustering and WHOIS-related signals.
(3) Explore whether certain technical profiles lead to systemic overblocking, indicating potential biases or lack of precision in PCS.

In doing so, our research aims to contribute to a better understanding of how PCS operate beneath the surface, moving beyond content analysis toward a nuanced view of how technical infrastructure shapes online access. We also seek to provide practical insights for improving the transparency and fairness of future parental control technologies.

3 RELATED WORK

Parental control systems are widely used to protect children from inappropriate online content. These systems typically rely on filtering mechanisms that are assumed to be based on the content of websites. However, recent research shows that technical features of websites, such as IP addresses, domain registrars, or DNS configurations, may also influence whether a site is blocked. This section reviews related studies and explains how they inform our investigation.

In one of the most detailed studies on PCS, Duchaussoy [3] looked at how commercial parental control tools work. He found that many PCS rely on external classification databases and DNS-based filtering. These tools often use third-party services to decide if a site is safe. Duchaussoy's work showed that PCS systems do not just look at website content but also at how a site is categorized in external databases. However, his study did not explain which specific technical domain features (like IP or WHOIS data) influence blocking.

The infrastructural dependency of filtering systems became further evident in studies focused on DNS manipulation and DNS over HTTPS (DoH). Trevisan et al. [15] presented methods for automatically detecting DNS manipulations, demonstrating how DNS-based interventions can influence access control systems like PCS. Similarly, Borgolte et al. [2] explored how DNS over HTTPS (DoH) affects filtering on the internet, especially in parental control and malware protection contexts. They discovered that encryption and DNS resolver choices could influence what is blocked. Although their study focused on network-level changes, it supports the idea that infrastructure and technical setup can drive filtering outcomes.

Hynek et al. [4] expanded on this by exploring how DoH protocols could be abused, and found that technical layers like DNS encryption and routing affected how websites were filtered. Their findings suggested that PCS could misuse encrypted protocols to enforce broader or hidden filtering rules. They raised concerns about transparency but did not perform a statistical analysis of technical features at the domain level. But highlighted the tight link between DNS infrastructure and filtering behavior.

Beyond DNS-level filtering, other research has looked at how PCS manage traffic at the home network level. Kamarudin et al. [7] developed PiWall, a home traffic controller that enables parental control and monitoring through localized DNS and IP filtering. This shows that even end-user controlled filtering solutions rely heavily on technical network-layer features rather than just website content.

From a broader perspective on filtering mechanisms, Magnusson [9] provided a broad review of DNS filtering methods across the internet. He concluded that many systems rely on superficial technical attributes like top-level domains or registrars, especially in large-scale DNS filters. While his study was theoretical, it highlighted the need for concrete evidence, which our research aims to provide. Similarly, Spaulding [13] developed D-FENS, a DNS-based filtering system to detect malicious domains. He found that filtering was sometimes based on domain history and reputation rather than content. Although this work focused on security systems, not parental controls, the mechanisms are similar and point to the same issue: technical features can lead to false positives in blocking.

Even though many studies have looked at PCS filtering, there is still an important gap. So far, no research has done a large, cross-system, statistical study to see how specific technical features—like WHOIS completeness, registrar names, or IP address ranges—are linked to the chance that a website gets blocked. Most earlier work focused on single PCS systems, general DNS behavior, or mistakes in content categories. None of them treated technical profiling as a separate reason why sites get blocked. This gap in the literature is what our study aims to address.

4 METHODS OF RESEARCH

4.1 WHOIS Data

Before explaining how we used WHOIS data in our analysis, it is important to explain what WHOIS actually is and why it matters in the context of this research.

WHOIS is a public database that stores registration information about internet domains. Whenever someone registers a domain name (like example.com), they provide details such as the registrar, creation date, expiration date, domain owner contact information, and name servers. This data is collected and made available through WHOIS services. The main purpose of WHOIS is to give transparency over who owns and manages a domain name.

Over time, privacy regulations like the GDPR have reduced the amount of information that is publicly visible in WHOIS records. However, many technical fields, like the registrar name, domain age, and name server provider, are often still available. These fields can give important hints about the background of a website—such as whether it was recently registered, who manages it, and where it is hosted. [5]

For our study, WHOIS data is especially relevant because by analyzing WHOIS attributes, we aim to find out whether such technical factors increase the likelihood of blocking.

4.2 Dataset

The project will use a given dataset with the following files. One of the files is a Tranco list containing a million domains—these domains constitute the input to the parental control systems. Each domain in the Tranco list has been mapped to one or more content categories using the Cisco Umbrella Investigate API. These categories allow us to measure blocking frequency across different types of content. This same Tranco list and its classification methodology were employed by Liberato et al. in their large-scale evaluation of parental control systems [8].

The second type of files are the blocklists generated by each of the parental control systems analyzed in the original study [8]. These systems include three router-based solutions (TP-Link, Netgear, and ASUS), two DNS filtering services (OpenDNS FamilyShield and DNS0.eu Kids), and one software-based tool (Norton Family). The blocklists correspond directly to the measurement data used in that study, indicating which domains were blocked versus which remained accessible. For our analysis, we compared these blocklists against the Tranco Top 1 Million list.

This study aims to analyze and compare the technical attributes of domains that are blocked and not blocked by Parental Control Systems (PCS). The methodology involves five main stages: feature extraction, data enrichment, analysis, and interpretation.

Each domain in the Tranco list will have technical features extracted using a combination of DNS lookups and WHOIS queries. These features aim to provide insight into whether certain technical characteristics correlate with blocking behavior by Parental Control Systems (PCS). Due to privacy laws as mentioned in the previous sections, rate limits, and variable data availability across domain registrars, it can occur that the requested features we plan to extract are not available. The features we plan to extract are:

- **IP-addresses:** This is the most common and easiest feature to extract since it is public information and retrievable via DNS lookups. DNS resolution is significantly faster and more reliable than WHOIS queries, making this feature well-suited for large-scale extraction.
- **Registrar:** The registrar field, if available, will be retrieved from WHOIS records using the python-whois library. This feature may not always be accessible due to registrar policies or privacy restrictions.
- **Domain age:** Using the creation date obtained from WHOIS, we will compute domain age. This will allow us to analyze whether newly registered or older domains are more likely to be blocked.
- **Name servers and contact emails:** These additional WHOIS fields will be extracted where available, providing further metadata for identifying infrastructure patterns or correlations with block decisions.

The extraction pipeline will attempt to retrieve each of these features for all domains in the Tranco list. In the event of missing data or limited coverage, adjustments in scope or focus will be explained in the results section.

4.3 Data Analysis

After extracting the technical features, we plan to conduct a series of statistical and descriptive analyses to evaluate their relationship with domain blocking by PCS. Each technical feature will be tested independently to assess whether it is significantly associated with blocking. In addition, we will explore interactions between multiple features to determine whether certain combinations are more predictive of blocking decisions.

Domain category information will also play a central role in our analysis. We aim to evaluate whether some categories are blocked more frequently than others, and whether specific technical attributes have stronger associations with blocking within particular categories.

Our analysis plan consists of the following steps:

- **Descriptive Statistics:** We will begin by summarizing the distribution of each technical feature. This includes computing frequencies, averages, and medians for attributes such as domain age, and visualizing category distributions across blocked and non-blocked domains.
- **Comparative Analysis:** We will compare the prevalence of each technical feature between blocked and non-blocked domains. These comparisons will be supported by tables and plots to highlight patterns of over- or under-representation.
- **Statistical Testing:** We will use statistical tests such as Fisher's exact test and chi-square tests to formally evaluate whether observed differences are statistically significant. Odds ratios will be calculated to quantify the strength of associations.

This step-by-step approach is designed to provide a robust and interpretable understanding of how technical characteristics influence the likelihood of a domain being blocked by different parental control systems.

5 RESULTS

This section presents the results of our analysis on the technical attributes of domains blocked by Parental Control Systems (PCS). Building on the classification framework established by Liberato et al. [8], we use the same Tranco-based dataset and Cisco Umbrella category mappings. However, rather than focusing on the qualitative behavior of blocking policies, our aim is to identify whether technical domain features—such as IP subnets, WHOIS metadata, and DNS characteristics—correlate with blocking decisions. Our results add an infrastructure-focused dimension to the evaluation of parental control systems, offering new insight into how technical features may influence domain-level censorship.

5.1 WHOIS Field Coverage and Limitations

Before conducting any statistical tests, we first analyzed the availability of WHOIS data across our dataset. This step is important because missing data could distort statistical conclusions later on. Due to time and resource constraints (approximately two weeks runtime for full enrichment), we limited WHOIS enrichment to four high-risk categories as defined by Liberato et al. [8]: Adult Content, Gambling, Hate/Discrimination, and Terrorism. These categories were selected because they are the most likely to trigger blocking by PCS.

- **Adult Content:** 33,368 domains
- **Gambling:** 20,846 domains
- **Hate/Discrimination:** 152 domains
- **Terrorism:** 11 domains

As shown in Figure 1, IP addresses were present in more than 95% of domains, while fields like registrar, creation date, and contact email had much lower coverage. These differences must be considered when interpreting statistical tests later, as missing data can affect significance results.

To ensure that IP-related analyses were robust, we used DNS lookups to retrieve IP addresses for the full Tranco Top 1M list. This guaranteed broad and statistically stable results for IP subnet analyses. In contrast, tests based on WHOIS attributes (like registrar and name server) were conducted only on the WHOIS-enriched subset within the high-risk categories mentioned above. This separation ensures that each analysis uses the most complete and reliable data available.

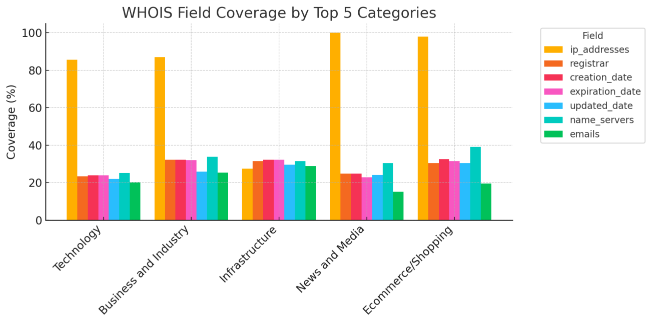

Figure 1: Percentage of domains with non-empty WHOIS fields by category

5.2 Category Distribution of Blocked Domains

Before testing technical features, we first checked if blocking behavior was still primarily driven by content categories. To do this, we compared the distribution of content categories of domains that were blocked, and domains that were not blocked. This acts as a baseline check to confirm that PCS are indeed still filtering based on content type.

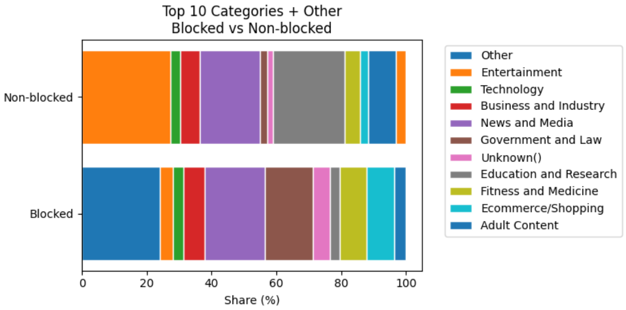

Figure 2: category distribution of non-blocked and blocked list of domains.

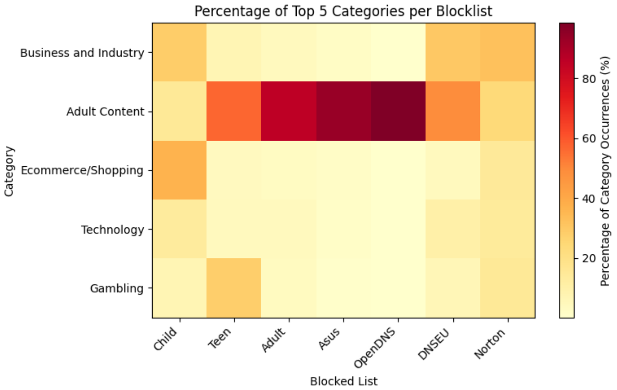

Figure 3: Heatmap of the top 5 most common categories in blocked domains across all PCS

Figures 2 and 3 show that sensitive categories like *Adult Content* and *Gambling* are overrepresented in blocked domains. This confirms that PCS still prioritize specific categories for filtering, reinforcing the need to control for category effects in our technical analyses that follow.

Also, a notable limitation in the current dataset is the relatively high number of domains labeled as **"Unknown"** in the Cisco Umbrella taxonomy . This labeling gap introduces classification noise, especially when trying to assess category-specific blocking trends. The presence of a large "Unknown" group can obscure true patterns, particularly if PCS apply blanket blocking rules to domains lacking clear categorization.

5.3 Subnet-Level Analysis

To investigate whether blocked domains cluster within specific parts of the IP address space, we conducted a subnet-level analysis. Specifically, we extracted /24 subnets from both blocked and non-blocked domains and ranked them by frequency.

The purpose of this analysis was to test for **infrastructure clustering**—whether certain IP ranges (subnets) are overrepresented in blocked domains. This could indicate that PCS sometimes block entire IP ranges instead of specific domains.

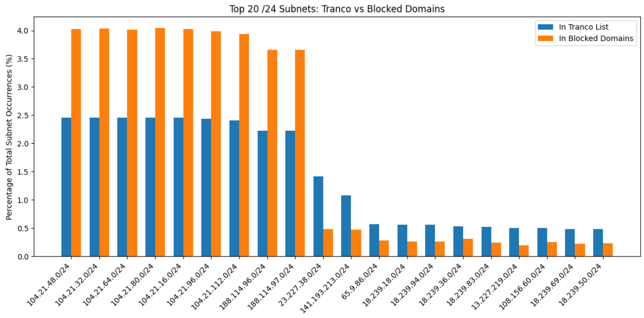

Figure 4: Top 20 most common subnets: comparison between blocked and non-blocked domains

Blocklist	Overlap Count
Child	39
Teen	42
Adult	38
Asus	34
OpenDNS	33
DNSEU	34
Norton	37

Table 1: Number of /24 subnets among top 50 of Tranco and each blocklist

Table 1 shows that a large portion of the top subnets in blocked lists also appear among the most common Tranco subnets. This suggests that PCS are blocking many domains hosted on shared infrastructure, whether intentionally or as collateral damage.

5.4 Category Alignment within Subnets

Following the subnet-level analysis, an important next step was to investigate whether the subnets that appeared frequently in blocklists showed any unusual concentration of specific content categories. In other words, we wanted to determine if the observed subnet clustering was simply the result of these subnets containing more “risky” or sensitive content, or whether the blocking behavior occurred even when the category composition of the subnet was comparable to the general internet.

To test this, we conducted a category alignment analysis between the top 100 most frequently blocked subnets from each PCS and the full Tranco Top 1 Million list, which served as a representative baseline for the general internet. For each comparison, we calculated the Pearson correlation coefficient (r) between the category distribution of domains within the blocked subnets and the category distribution in the overall Tranco dataset.

The Pearson correlation coefficient is a widely used statistical measure that captures the strength and direction of the linear relationship between two variables. In our case, each variable represents the relative frequency of each content category within a given set of subnets (for example, the blocked subnets versus the full Tranco list). An r-value can range from -1 to $+1$. A value close to $+1$ indicates a very strong positive correlation, meaning the two category distributions are highly similar in shape and proportion. An r-value near zero would suggest no correlation, while an r-value near -1 would indicate an inverse relationship. By applying this metric, we could objectively assess how closely the category mix of blocked subnets resembled the general internet baseline.

The results of this analysis revealed a remarkably high degree of similarity. For all PCS examined, the Pearson correlation coefficients exceeded 0.9, indicating a very strong positive relationship between the two distributions. This means that, despite being heavily targeted by blocking mechanisms, the content mix of these subnets remained closely aligned with the general distribution of website categories found across the broader internet.

This high correlation is statistically significant and highly unlikely to be the result of random chance. If PCS were deliberately targeting subnets because they contained disproportionate amounts of high-risk content—such as adult sites, gambling platforms, or hate speech—we would expect to see much lower correlation values. Such a targeted blocking pattern would manifest as a notable divergence in category composition between blocked subnets and the overall internet baseline. However, the consistently strong alignment observed across all PCS suggests otherwise.

The findings from this category alignment test provide strong evidence that PCS are not selectively blocking subnets based on their content profiles. Instead, it appears that infrastructural or technical factors—such as shared hosting environments or IP reputation heuristics—are driving these blocking decisions. The fact that entire subnets with a category composition typical of the general internet are still being blocked points towards an infrastructural bias within PCS filtering logic.

5.5 Statistical Tests: Field Presence vs. Blocking

To evaluate whether the **presence or absence** of specific WHOIS fields influences the likelihood of a domain being blocked, we conducted a series of statistical tests grounded in categorical data analysis. Given the nature of our dataset—where some WHOIS fields are sparsely populated and certain content categories contain relatively few blocked domains—we selected **Fisher’s exact test** as our primary method for assessing association.

Fisher’s exact test is particularly well-suited for this context because it evaluates the non-random association between two categorical variables within a 2×2 contingency table, without relying on large-sample approximations. This makes it robust even when cell counts are low, avoiding the limitations of other tests like the Chi-square test, which assume larger expected frequencies in each cell.

For each combination of PCS, WHOIS field, and domain content category, we constructed a contingency table contrasting the number of **blocked** and **non-blocked** domains with the **field present** versus **field absent**. Formally, each table was organized as follows:

We then applied Fisher’s exact test to each table to evaluate the null hypothesis that the presence of a given WHOIS field is *independent* of the domain’s blocking status. A **p-value below 0.05** was interpreted as evidence of a statistically significant association.

Given the potential for category-based confounding, we stratified our analysis by domain content category. This stratification was crucial to control for the fact that certain domain types (e.g., adult content, gambling, etc.) might inherently attract more blocking regardless of WHOIS field presence.

To summarize these results across multiple categories and PCS, we calculated **odds ratios (ORs)** and visualized the findings using **$\log_2$-transformed ORs**, as shown in Figure 5. An odds ratio greater than 1 indicates that the presence of a particular WHOIS field is associated with an increased likelihood of blocking, while an odds ratio less than 1 suggests a protective or negative association. Statistically significant associations, as determined by Fisher’s test, are marked with an asterisk (*) in the figure.

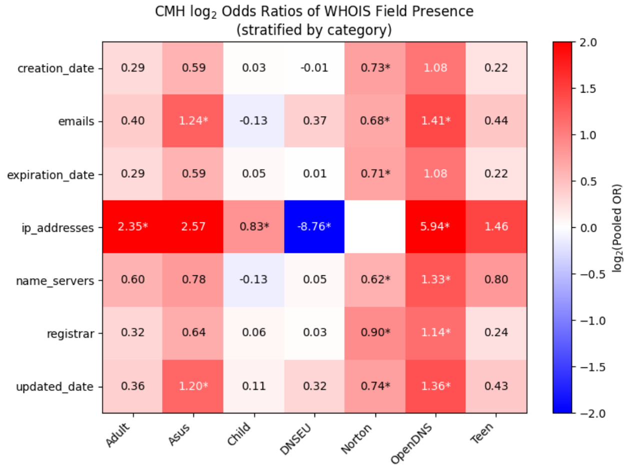

Figure 5: Figure 5: CMH Log$_2$ odds ratios of WHOIS field presence (stratified by category) “$*$” indicates $p < 0.05$.

Overall, our findings suggest that the presence of certain WHOIS fields—correlates with a higher likelihood of a domain being blocked. This pattern implies that Parental Control Systems may leverage WHOIS field visibility as a heuristic signal in their classification processes, or alternatively, that domains with more complete WHOIS records are more easily identified and subsequently targeted for blocking.

5.6 Value-Specific Blocking Behavior

Finally, we tested whether **specific values** within WHOIS fields (e.g., particular registrars or creation years) are linked to blocking likelihood.

For this, we applied the **Cochran-Mantel-Haenszel (CMH) test**, which allows for pooled odds ratio calculations across stratified data (in our case, stratified by content category). This method helps control for category-based effects and focuses on the impact of each specific WHOIS value.

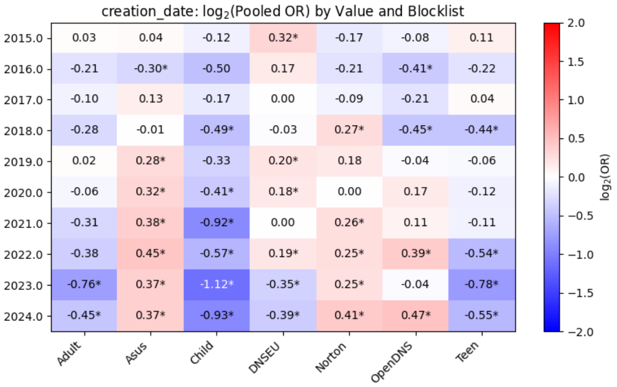

Figure 6: Log$_2$ pooled OR for top creation years by PCS. "$*$" indicates $p < 0.05$.

Key findings from our analysis reveal several WHOIS value patterns that appear to influence PCS blocking behavior. One of the most striking observations is the elevated block rate for domains created in 2023 and 2024 (Figure 6). This trend aligns with findings from Unit 42 and Akamai, both of which report that newly registered domains (NRDs) are heavily targeted in cybersecurity due to their frequent use in phishing, malware distribution, and other malicious activities [14, 16]. For example, Unit 42 found that more than 70% of NRDs were classified as malicious or suspicious, while Akamai detected millions of newly observed harmful domains within their first weeks online. Given this broader threat landscape, it is plausible that PCS proactively block newer domains as a risk mitigation strategy.

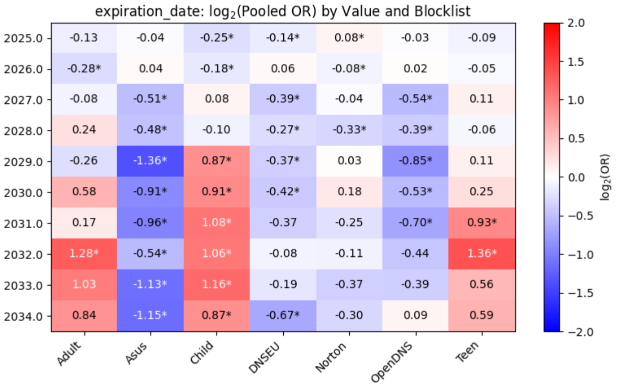

Figure 7: Log$_2$ pooled OR for top expiration years by PCS. "$*$" indicates $p < 0.05$.

Further supporting this interpretation, we observed that domains that are due to expire soon—such as those expiring in 2024—also exhibited higher block rates (Figure 7). This corresponds with research from WhoAPI and recent DNS abuse studies, which highlight that many malicious domains are registered for short durations to support time-limited attacks like phishing or spam [17, 18]. PCS systems may therefore interpret short expiration periods as a red flag for domain trustworthiness.

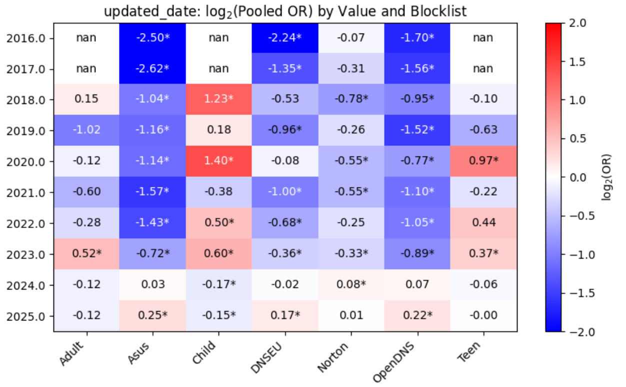

Figure 8: Log$_2$ pooled OR for top updated years by PCS. "$*$" indicates $p < 0.05$.

Another notable pattern concerns the last update year of domains. Domains that were recently updated showed a significantly higher likelihood of being blocked (Figure 8). This is consistent with industry findings that malicious actors frequently modify WHOIS information or update domain infrastructure to evade detection and blacklists [16, 18]. Such frequent changes may signal instability or suspicious activity to PCS filters.

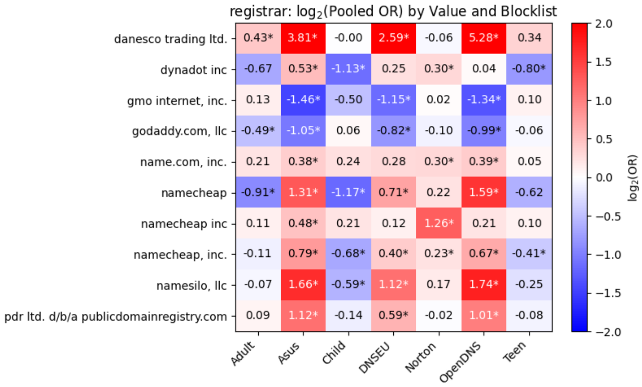

Figure 9: Log$_2$ pooled OR for registrars by PCS. "$*$" indicates $p < 0.05$.

Looking at infrastructure-related fields, we found that domains using large DNS hosting providers—such as Cloudflare and Namecheap—appeared more frequently in blocked lists (Figure 10,9). While these providers are widely used across the internet, the elevated block rates we observed remain significant even after controlling for domain category. This suggests that PCS may be applying registrar-based heuristics rather than simply reflecting registrar popularity. This observation reflects broader industry concerns, as highlighted by Unit 42, which noted that attackers often prefer large, scalable providers that enable rapid deployment and easy registration processes [16]. While these providers themselves are not inherently risky, their widespread availability makes them popular among malicious actors.

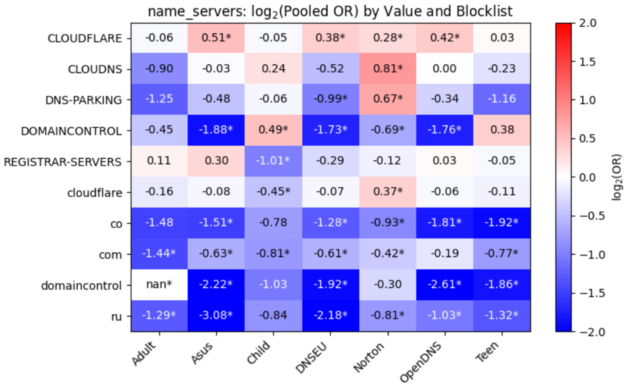

Figure 10: Log$_2$ pooled OR for top name-server vendors by PCS. "$*$" indicates $p < 0.05$.

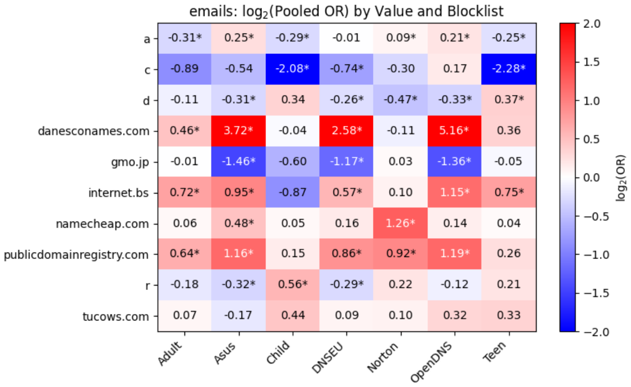

Figure 11: Log$_2$ pooled OR for top email providers by PCS. "$*$" indicates $p < 0.05$.

Finally, domains with contact emails from free, anonymous email services like Gmail and Yahoo also experienced elevated block rates (Figure 11). According to WhoAPI, the use of such email providers in WHOIS records is commonly associated with domains intended for fraudulent or malicious purposes [18].

Together, these findings suggest that PCS systems do not solely rely on the presence of WHOIS fields but actively consider the specific values within these fields when making blocking decisions.

5.7 Future Work

Our study showed that DNS reachability and WHOIS metadata can influence parental control decisions, but there is still much to cover. A key next step is to enrich WHOIS information for every domain in the Tranco list, not just a handful of high-risk categories: bulk RDAP feeds or commercial APIs could give us reliable registration dates and contact details and remove any sampling bias. At the same time, we should improve how we label and group domains in the Tranco list itself. Rather than relying solely on the static Umbrella categories, we could apply automated content analysis or machine-learning classifiers to capture emerging clusters (for example, distinguishing "streaming video" from generic "entertainment") and reduce the number of "Unknown" entries. Incorporating hierarchical or multilabel taxonomies - where a site can belong to subcategories like "News Finance" or "Social Messaging" - would let us see more nuanced blocking patterns. On the network side, going beyond the /24 CIDR counts by using reverse DNS lookups, TLS certificate fingerprints, or CDN provider identification might reveal provider-specific biases in blocklists. Finally, adding geolocation and autonomous-system data and pairing our measurements with longitudinal scans or user experience studies would show how policies evolve over time and what real impact they have on everyday browsing. By broadening our WHOIS coverage and our domain categorization, we can gain a richer, more accurate picture of how technical signals shape parental control filtering.

6 CONCLUSION

Looking closely at the results, we can see how Parental Control Systems (PCS) decide to block websites—not just based on what kind of content is on the site, but also based on technical details. Although the type of website (category) still plays a big role, we found that many PCS also use technical clues to decide what to block.

One of the clearest results we found was that blocked websites often come from the same small groups of IP addresses (called /24 subnets). These subnets are shared by many websites and show up in the blocklists of several PCS. This means that PCS might be blocking groups of websites together because they are hosted on the same servers, even if only some of them are harmful.

We also saw that certain technical details in WHOIS data are linked to blocking. Websites with active IP addresses, common email providers (like Gmail or Yahoo), and known name servers are more likely to be blocked. Even though these details don't tell us directly about the content, they might be used as warning signs by PCS.

In addition, we found that some domain registrars (like GoDaddy and Namecheap) are seen more often among blocked sites. This could be because these companies are popular for hosting websites with risky content—or because PCS are more likely to block sites from certain registrars. We also noticed that newer websites (especially from 2023–2024) are blocked more often. This shows a bias against recently created websites, which are frequently used in phishing and malware campaigns, as documented in cybersecurity reports by Unit 42 [16] and Akamai [14].

Another important finding was that the types of websites blocked within shared IP groups closely resemble the overall distribution of website types found across the internet. This suggests that PCS are not always carefully picking which sites to block—they may just block entire chunks of the internet, including safe websites, based on technical infrastructure signals rather than content indicators.

We also observed that not all PCS behave the same way. Some systems, like ASUS and OpenDNS, showed stronger blocking patterns for newly registered or recently updated domains, suggesting a more aggressive focus on technical freshness signals, which aligns with known cybersecurity heuristics for identifying risky domains [14, 16]. Others, like TP-Link and DNSEU, showed weaker correlations with these features, indicating that they may prioritize different heuristics or place more weight on category-based filtering. Similarly, block rates for domains hosted on large infrastructure providers like Cloudflare and GoDaddy were noticeably higher in Norton Family and OpenDNS but less pronounced in other PCS. These differences suggest that while technical profiling is a shared trend, its implementation is far from uniform across vendors.

Overall, this study shows that PCS use both content-based and technical signals to block websites. This mix of methods can lead to overblocking—where safe websites are blocked just because of how or where they are hosted. Since PCS don't clearly explain how they make these decisions, users may not understand why certain sites are blocked, and it's hard to hold PCS companies accountable for mistakes.

7 AI STATEMENT

This research project was assisted by AI (OpenAI o4-mini). The AI helped correct grammar, ensure proper academic formatting, and suggest relevant sources. All references cited are real; their bibliographic entries were formatted by the AI based on the provided sources. Additionally, we consulted the AI to better understand the behavior and limitations of the `python-whois` library—clarifying which date fields are registry defaults versus true domain registration dates, and seeing what data could be used from the requests.

REFERENCES

[1] Blocksi. Keywords your school's content filter should block and why, 2023. Accessed: 27 June 2025.

[2] Kevin Borgolte, Tanmay Chattopadhyay, and Nick Feamster. How dns over https is reshaping privacy, performance, and policy in the internet ecosystem. *Princeton University Working Paper*, 2019.

[3] Quentin Duchaussoy. *Security and Privacy Analysis of Parental Control Solutions*. PhD thesis, Concordia University, 2020.

[4] Karel Hynek, Dmitry Vekshin, Jan Luxemburk, and Tomas Cejka. Summary of dns over https abuse. In *2022 IEEE Symposium on Security and Privacy Workshops (SPW)*, pages 104–109. IEEE, 2022.

[5] ICANN. Temporary specification for gtld registration data, 2018. Accessed: 27 June 2025.

[6] Internet Matters. Children's wellbeing in a digital world: Index report 2024, 2024. Accessed: 27 June 2025.

[7] Siti Kamarudin, Nur Izzati Abd Razak, and Mohd Izzuddin M. Shuhud. Piwall as a home traffic controller: Enabling parental control and monitoring. *Bulletin of Electrical Engineering and Informatics*, 13(2):856–865, 2024.

[8] Matteo Liberato, Antonia Affinito, Bernd Meijerink, Mattijs Jonker, Alessio Botta, and Anna Sperotto. To block or not to block? evaluating parental controls across routers, dns services, and software. In *Proceedings of the IFIP Network Traffic Measurement and Analysis Conference (TMA) 2025*, page –. IFIP, 2025.

[9] Johan Magnusson. Survey and analysis of dns filtering components. *arXiv preprint arXiv:2401.03864*, 2024.

[10] National Coalition Against Censorship. Internet filters, 2024. Accessed: 27 June 2025.

[11] Ofcom. Children and parents: Media use and attitudes report 2024, 2024. Accessed: 27 June 2025.

[12] Pew Research Center. How teens and parents approach screen time, 2024. Accessed: 27 June 2025.

[13] Justin Spaulding. *D-FENS: DNS Filtering and Extraction Network System for Malicious Domain Names*. PhD thesis, University of Central Florida, 2018.

[14] Akamai Technologies. Newly observed domains: Discovered 13 million malicious domains, 2022.

[15] Martino Trevisan, Idilio Drago, and Marco Mellia. Automatic detection of dns manipulations. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 4736–4745. IEEE, 2017.

[16] Palo Alto Networks Unit 42. Newly registered domains: Malicious abuse by bad actors, 2022.

[17] Xinjian Wang, Rohan Sahay, and Md Sazzadur Rahman Sarker. Registration, detection, and deregistration: Analyzing dns abuse for phishing attacks, 2024. arXiv preprint arXiv:2502.09549.

[18] WhoAPI. Why are newly registered domains important for cybersecurity?, 2023.