

Hidden Secrets in the arXiv: Discovering, Analyzing, and Preventing Unintentional Information Disclosure in Source Files of Scientific Preprints

Jan Pennekamp^{*,§}, Johannes Lohmöller^{*,§}, David Schütte^{*}, Joscha Loos^{*}, Martin Henze^{‡,¶}

^{*}Communication and Distributed Systems, RWTH Aachen University, Germany · {lastname}@comsys.rwth-aachen.de

[‡]Security and Privacy in Industrial Cooperation, RWTH Aachen University, Germany · {lastname}@spice.rwth-aachen.de

Abstract—Preprints are essential for the timely and open dissemination of research. arXiv, the most widely used preprint service, takes the idea of open science one step further by not only publishing the actual preprints but also \LaTeX sources and other files used to create them. As known from other contexts, such as GitHub repositories, and anecdotally exemplified for arXiv, making source code publicly available risks disclosing otherwise “hidden” information. Consequently, the public availability of paper sources raises the question of how much sensitive content is (unintentionally) disclosed through them.

In this paper, we systematically answer this question for all 2.7M arXiv submissions with available source files across three dimensions of source file-induced information disclosure: ① inclusion of unnecessary files, ② metadata embedded in files, and ③ irrelevant content in files such as source code comments. Our analysis reveals that nearly every arXiv submission contains some form of “hidden” information. Notable findings range from links to editable web documents for internal coordination over API and private keys to complete Git histories.

While different tools promise to remove such information from source files, we show that they fail to reliably achieve the intended cleaning functionality. To mitigate this situation, we provide ALC-NG to comprehensively remove files, metadata, and comments that are not needed to compile a \LaTeX paper.

Index Terms—information disclosure; source code; \LaTeX ; arXiv

1. Introduction

Preprints, i.e., early versions of academic papers already made available before formal peer review [1], are an indispensable tool for the timely and open dissemination of academic research, with first research fields such as machine learning even shifting to a preprint-first “publication” model [2]. Since its inception in 1991 [1], arXiv [3] has established itself as the largest and most frequently used preprint service [4] and is widely accepted in disciplines such as computer science, physics, and mathematics, with >20 000 preprints published every month [5]. In addition to the PDF of the preprint, arXiv also publishes the source code of all papers written using \LaTeX ($\approx 93\%$ of submissions). While initially intended to enable the reuse of notations and provide learning opportunities [6], publishing preprint source

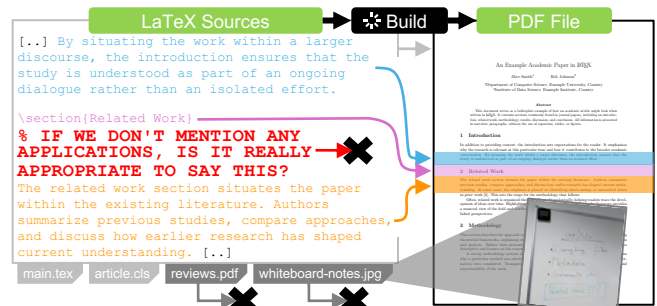


Figure 1. \LaTeX sources might contain comments or dangling files that are irrelevant to or unused in the compiled PDF, respectively. When being distributed regardless, e.g., via arXiv, they disclose (sensitive) information.

code enables research such as citation prediction [7], studying inconsistencies in documents produced using \LaTeX [8], or reconstructing \LaTeX sources from PDFs [9], [10].

However, besides providing many benefits [11], making the source code of academic papers publicly available also introduces substantial information disclosure risks [6]. As exemplarily highlighted in Fig. 1, the two main risks stem from (i) irrelevant content, e.g., in the comment sections of \LaTeX files, and (ii) unused files within the uploaded archive that are not needed to build the paper. Notably, these risks come on top of information disclosure through metadata embedded in the papers themselves, which is well known and researched [12]–[18]. To the larger public, the risks specific to publishing sources of preprints were first anecdotally mentioned on Twitter/X [19]. However, the underlying issue gained wider attention when a paper on GPT-4 purposefully omitted details concerning its toxicity probability and how it complies with intellectual property laws [20].

Still, despite arXiv maintaining its own build pipeline and warning about corresponding issues [6], proactive mitigation of those issues at scale seems to be lacking. While arXiv endorses [6] Google’s arxiv_latex_cleaner (ALC) [21], [22]—a tool claiming to “easily clean the \LaTeX code of [...] paper[s]”—it does not enforce any actions on the submitting authors. Similarly, while anecdotal evidence and a few academic works [23], [24] hint at the issue of unintentionally distributing comments and other data through \LaTeX source files, a detailed analysis of the prevalence and scale of security problems associated with the distribution of papers’ sources is lacking. However, the lack of understanding these

§. Equal contribution. ¶. Also with Fraunhofer FKIE.

problems challenges us in assessing respective risks, notifying affected users, and taking informed countermeasures.

Tackling this research gap, we set out to quantify the prevalence of “hidden” information, including sensitive data, within *all* \LaTeX sources published by arXiv. In particular, we consider the availability of unused (dangling) files, embedded metadata, and compilation-irrelevant content (e.g., comments in \LaTeX files) to get a holistic picture. Indeed, our findings show that “hidden” information within arXiv is a widespread issue: 88 % of papers with \LaTeX sources are affected by at least one type of “hidden” information, making information available for everyone to “discover” and potentially exploit. Our analysis reveals substantial amounts of content that were either not meant for distribution or should not be publicly distributed, including reviewer comments, author discussions, survey data with personally identifiable information (PII), as well as links to internal documents and reports without any restrictive access control. Perhaps surprisingly, papers related to top-tier security conferences contain significantly more “hidden” information, despite this community’s expectable awareness of such issues. Overall, the prevalence of issues concerning sensitive information prompted us to initiate two responsible disclosure campaigns (cf. Appendix A).

Contributions. To provide a profound understanding of the prevalence and severity of “hidden” information within arXiv submissions, better assess corresponding risks, and lay the foundation for mitigations, our main contributions are:

- We present a methodology to reliably detect irrelevant content in \LaTeX sources and identify unused files. Unlike current mitigation approaches that solely rely on heuristics, our approach utilizes an abstract syntax tree and `pdflatex`’s recorder option for precise identification of potentially critical “hidden” information (Section 4).
- We conduct the first large-scale analysis of “hidden” information within preprint sources based on a comprehensive dataset of \LaTeX sources at arXiv (Section 5), covering *all* 2.7M submissions with source files.
- We reveal that the majority (88 %) of papers feature (unnecessary) information within their arXiv submission, with striking findings such as secrets (API tokens, private keys, passwords), sensitive content (links to internal documents, disclosure of FTP and SSH servers), and uncommented text fragments and discussions among coauthors (Section 6).
- To address the identified serious and prevalent issues, we develop and release ALC-NG—an improved \LaTeX cleaner that reliably prevents “hidden” information distribution, mitigating 90 % of “hidden” information issues (Section 7).

Open Science Statement. We are committed to open science and thus publish aggregated data of our evaluation [25]. We also release our \LaTeX source file cleaner—ALC-NG—for public use [26]. In compliance with our ethical considerations (Appendix A), we refrain from publishing our analysis code.

2. Preprint Repositories and arXiv.org

In academia, a preprint corresponds to a scientific paper that is distributed prior to peer review or official publication [1], [4], [27]. While they can be shared any time,

around 40 % of preprints get (finally) formally published [4]. Independent projects (see DOAPR [28], the directory of open access preprint repositories, for an overview) collect preprints on specific domains, e.g., bioRxiv and medRxiv, or for specific purposes, establishing different preprint repositories. Some, such as SSRN or TechRxiv, are even operated and/or endorsed by academic publishers. In 2021, Xie et al. [4] studied the use of preprint repositories, particularly noting the popularity of arXiv [3] and identifying a substantial increase in usage in computer science over the past few years.

Established in 1991 to consolidate the then prevalent preprint sharing via email in physics research [1], arXiv [3] has evolved into a free, open-access preprint repository, hosting more than 2.9M research papers [5] of different domains [29], including natural sciences, computer science, quantitative finance, electrical engineering, and economics. The importance of preprints has been established decades ago [30], [31], and the observed growth is domain-independent, i.e., it also holds for computer security [32]. Since its inception, both the number of submissions and downloads have been increasing continuously [5], [29], [33].

The general submission process at arXiv works as follows. After writing the paper, an author creates a submission by uploading the \LaTeX sources of the paper. arXiv then compiles the sources and conducts some basic checks to verify whether the paper is eligible for inclusion. This process is not a proper peer review, but rather a means for maintaining quality within the repository [29], [34]. Upon successful vetting, the paper is publicly announced, and a *permanent, versioned record* is created. After the initial announcement, authors can decide to upload a new, improved (corrected) *version* that amends the original record or withdraw a version for all kinds of reasons [35]. In the latter case, arXiv still retains the original submission but adds a dedicated notice of withdrawal.

For different reasons [6], including accessibility and improved long-term maintainability, arXiv requires authors to submit \LaTeX sources if available. Accordingly, more than 90 % of submissions are distributed with their sources [11].

3. Related Work

Research analyzing arXiv has largely focused on non-security aspects, such as (i) studying how academic papers are written using \LaTeX [23], [24], (ii) understanding the impact of arXiv submissions [36] and the influence of preprints on citations [36]–[41], (iii) conducting analyses of authorship and publication forms [42]–[47], (iv) detecting plagiarism [48], text reuse [49], and AI-generated texts [50], (v) collecting reasons for withdrawing preprints [35], (vi) training and evaluating language models [51]–[53], and (vii) reconstructing \LaTeX sources [54], analyzing document layouts [55], or extracting content from \LaTeX -compiled PDFs [56].

However, despite enabling all these diverse research streams, publishing sources of \LaTeX documents inevitably carries the risk of unintended information disclosure [6]. Indeed, related work analyzing source code in GitHub repositories [57]–[61], Docker containers [62], and mobile apps [63] has identified a substantial risk of disclosing information such

as secret API keys, private keys, or passwords. Consequently, different tools to sanitize \LaTeX documents before publishing them on arXiv have been developed [6], [21], [64]–[66]. While we refer to Section 7.1 for a detailed survey of those tools and their capabilities in accurately sanitizing \LaTeX documents, anecdotal evidence of (remaining) critical content in sources of arXiv submissions [19], [20], [23], [24] suggests that these tools are either not widely applied or ineffective in removing (all) critical content.

Looking beyond source files, information leakage resulting from published documents such as PDF files [13], [16]–[18], images [67], [68], and other formats [12], [14], [15] has been analyzed, uncovering several common issues, including GPS locations where images were taken, author names, and other (partially) private information unintentionally stored primarily in metadata fields. To our understanding, these issues of “hidden” information in published documents are well-researched and well-understood, with practical tools for metadata removal being available [69]–[71]. Further work has focused on detecting malicious PDFs [72]–[74], which is only partly applicable to arXiv-published papers given its service-provided build pipeline, or the identification of privacy issues within PDF-based publication processes [75].

Consequently, research is well aware that “hidden” information in documents, including metadata, can be a threat as it may unintentionally reveal sensitive details, e.g., GPS locations in images or author details in rich text documents. At the same time, we have arXiv—the largest and most popular preprint service—which, in contrast to most other preprint services, requires and publicly distributes author-submitted \LaTeX source files. Despite selective attention on the Web [19], [20] and in academia [23], [24] as well as warnings issued by arXiv [6], research has not yet assessed the prevalence of this phenomenon at scale. Concurrent to our work, two preprints with specific foci—a LLM-driven security audit of source files [76] and assessing the content of nonrequired files within sources [77]—also study \LaTeX source files, albeit at significantly smaller scale. Still, they highlight and emphasize the relevance of our research.

4. Methodology For Analyzing \LaTeX Sources

The goal of our work is to discover, quantify, and analyze the distribution of “hidden” information at the popular preprint service arXiv. These findings also serve as a foundation for mitigation strategies to prevent the inclusion of “hidden” information. In line with research in other contexts (Section 3), we identify three key dimensions of “hidden” information in arXiv submissions: ① \LaTeX sources bundled with *dangling files*, ② embedded *metadata*, and ③ *content/text* irrelevant to the build process.

In the following, we present our methodology to comprehensively uncover “hidden” information across these three dimensions with high accuracy to lay the foundation for analyzing arXiv submissions and their source files at scale.

Dimensions of “Hidden” Information. Our methodology for discovering “hidden” information aligns with the three dimensions along which such information can occur.

① *File Level.* At the file level, we classify source files as either required or dangling files. *Required* files are strictly needed to compile the \LaTeX sources into a PDF document, while *dangling* files are not used during the build process and thus do not need to be part of the served sources.

We align our detection of dangling files with arXiv’s long-used submission system (version 1.0). That is, we mark all \LaTeX root files with `documentclass` or `documentstyle` as required and then recursively mark all additionally included files as required. We further support `00README(.XXX)` files [78], present in 8.8% of all arXiv submissions with sources, in which authors may optionally provide additional information to the submission system, e.g., to override arXiv’s default inclusion behavior for files.

② *Metadata Level.* Most file types further support the inclusion of metadata to add supplementary information and provenance data. Since related work (cf. Section 3) has shown that metadata may leak sensitive information [79], we rigorously analyze all source files using `exiftool` [69].

③ *Content/Text Level.* Besides content within required files that is needed to produce the PDF document, source files on arXiv can also contain *irrelevant* content with no influence on the compiled PDF. In this regard, \LaTeX source files (beyond other file types) are especially relevant given their intentional support for author-supplied comments [23]—a common practice in (academic) writing.

To identify irrelevant content in \LaTeX sources, our methodology relies on an abstract syntax tree representation of said \LaTeX document using a tree-sitter [80] grammar. Harnessing this structured representation of \LaTeX documents, we can reliably parse and extract irrelevant content, e.g., comments, by specifying the different ways irrelevant content can be embedded in \LaTeX sources. This approach fundamentally differs from the current state of the art in sanitization tools (cf. Section 7.1), which rely on heuristic approaches based on regexes. To ensure comprehensive coverage of the different concepts of embedding comments, we survey the treatment of irrelevant information by six state-of-the-art \LaTeX sanitization tools (cf. Section 7.1) and supplement these findings with our own experience from authoring more than one hundred academic papers with various research groups across three continents.

Syntax of Embedded Comments. Using this approach, we derive four widely-used and commonly-known concepts for embedding comments in \LaTeX documents that we employ to detect irrelevant content: ① standard \LaTeX line comments precluded by `%` or wrapped in a `comment` environment, ② any characters outside of \LaTeX ’s document environment, ③ non-taken branches of (custom) `if` statements, and ④ ignored arguments of `(re)newcommand`, `providecommand`, `Declare*`, or `def` constructions. Complementing the content-level dimension, we further consider all content in dangling files as irrelevant.

Given \LaTeX ’s expressiveness, we may theoretically miss other irrelevant content that does not match the syntax described above; for example, `csname` constructions to insert content or author-implemented conditionals. However, our investigation into the usage frequency of custom commands

did not reveal meaningful findings. Thus, such practices seem rare within arXiv sources, i.e., they are not used at scale.

Handling Multiple Versions of a Submission. When analyzing arXiv papers, multiple *versions* (i.e., updates) of a single paper must be considered (cf. Section 2). Issues present in one version may not be part of another version of the same paper (either the issue gets introduced in a later version or removed during an update). To avoid overcounting “hidden” information, we thus perform a worst-case aggregation over versions. That is, we consider an issue to be present for a given paper if we discover it in *at least* one version, not necessarily the most recent one. As the sources of older versions remain accessible for anyone, this approach is reasonable and accurately quantifies information leakage.

Validity. By relying on a structured parsing approach based on an abstract syntax tree (instead of relying on heuristics, e.g., based on regular expressions), our approach intentionally constitutes an accurate lower bound of detected irrelevant content. Furthermore, as our coverage of the different concepts of embedding comments relies on best practices from existing sanitization tools and extensive authoring experience, we are convinced that we capture the vast majority of irrelevant content. While \LaTeX , as a Turing-complete language, theoretically offers additional custom approaches for embedding comments, as part of our analysis, we did not discover such peculiarities at scale and thus estimate their prevalence to be negligible in practice. As such, our assessment methodology (a) reasonably minimizes false negatives (undetected irrelevant content), and (b) yields almost no false positives (relevant content falsely flagged as irrelevant). Alongside turning our comprehensive detection methodology into a next-generation sanitization approach (cf. Section 7.3), we will show that our abstract syntax tree-based approach outperforms all existing sanitization tools in simultaneously minimizing false negatives and false positives while not visually changing compiled documents.

To the best of our knowledge, we are the first to apply a methodological approach to assess the prevalence of “hidden” information in \LaTeX source files at scale. Our methodology ensures an assessment that minimizes false positives, i.e., we do not label content as irrelevant if it is strictly needed when compiling the PDF from source. In contrast, although covering all prevalent approaches to embed irrelevant content, our approach does not guarantee completeness; rare custom cases might be missed. Consequently, our results constitute a reliable lower bound of discovered “hidden” information.

5. Dataset Covering All arXiv \LaTeX Sources

For our analysis, we have curated a dataset covering the source files of *all* arXiv submissions from its launch in 1991 up to and including 12/2025. For 1991–10/2020, we retrieved the archive.org collection [81]. To retrieve the remaining months and fill any gaps, we relied on the official arXiv-bulk dataset access over Amazon S3 storage [82]. Since these sources only serve one version of a paper, we further sampled full version histories using the official arXiv API while complying with their rate limits (cf. Appendix A).




Figure 2. Annual volume of arXiv submissions with and without sources.

After retrieval, we validated the completeness of our dataset: (i) we iterated over the paper ID space to check for gaps (some IDs are not associated with a submission), (ii) we marked all withdrawn versions and papers (similar to WithdrawarXiv [35]), and (iii) we verified that submissions without source files in our dataset are indeed without sources. For the latter two steps, we were in contact with arXiv.

In total, our dataset covers 2.7M papers with source files, which corresponds to 93% of all papers on arXiv. Fig. 2 shows their distribution over the years, with 2013 representing the relative lower bound (91%) and 1991 the upper bound (100%). Across the papers with source files, we sampled 1.1M papers with more than one version with source files, accumulating to a total of 4.3M versions.

6. Analysis of arXiv Source Files at Scale

Our dataset of all 2.7M submissions with available source files published on arXiv (by the end of 2025) enables us to discover and analyze “hidden” information within sources made available by arXiv and thus derive a comprehensive picture of the prevalence of unintentional information disclosure in source files of scientific preprints. To this end, we first provide a high-level overview of our findings, summarizing frequent types of “hidden” information and reporting changes over the years, arXiv categories, and submitted paper versions (Section 6.1). Subsequently, we attribute our findings to the different dimensions of the problem (cf. Section 4) and sequentially discuss them in greater detail (Section 6.2).

We focus on assessing the prevalence of “hidden” information and grasping its sensitivity. Consequently, the implications of content-specific changes over time, e.g., the removal or addition of authors, the effects of changing templates, or the creation of citation graphs, which related work already covered partly (cf. Section 3)—albeit at a comparatively small scale—are not in our scope.

6.1. A Birds-Eye View at “Hidden” Information

We begin our analysis by examining the rate of “hidden” information at the file, metadata, and content/text levels.

Guiding Question. *How prevalent are the three dimensions of “hidden” information in source files of preprints?*

To provide insights into this guiding question, we filter for interesting content (Section 6.1.1), quantify frequent issues (Section 6.1.2), and analyze differences across submission categories, versions, and years (Section 6.1.3).

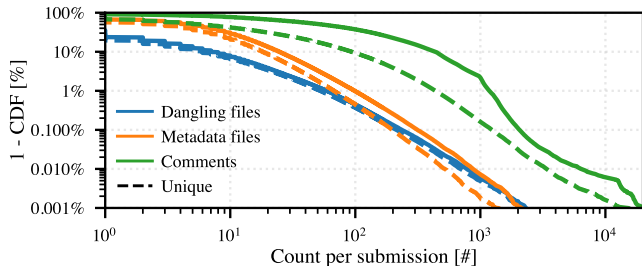


Figure 3. Complementary cumulative distribution functions (1-CDFs) of total and unique “hidden” information per arXiv submission across the three dimensions: dangling files, metadata key-value pairs, and \LaTeX comments.

6.1.1. Filtering For “Interesting” Content. After processing the entire dataset of arXiv sources, we obtain a massive volume of “hidden” information for analysis across all considered dimensions, including 12 M dangling files, 144 M metadata key-value pairs, and 644 M comments. Fig. 3 shows their distribution across submissions. However, a substantial fraction of this massive volume of “hidden” information repeats across (many) submissions (e.g., unused template files or comments in templates) and thus likely does not constitute potentially unintentionally-disclosed information.

“Interesting” Content. To instead focus only on potential cases of unintentional information disclosure, we define “interesting” content as any content that is (almost) *unique* to an individual submission, following the intuition that “hidden” information occurring only in one submission likely was added by the authors of this submission (and not inherited, e.g., from templates). More specifically, we filter out any dangling file, metadata value, or comment that we find in more than two submissions (to account for rare duplicates by chance or reuse by the same authors).

In Fig. 3, we visualize complementary cumulative distribution functions (1-CDFs), detailing the number of unique and total information for the three dimensions.

File Level. Out of 50 M total files, we find 45 M (91 %) to be unique (occurring at most twice) based on SHA256 fingerprints (44 M/88 % for strict uniqueness). This methodology filters out 181 k distinct files that we, on average, find in 23.6 different submissions (median=3, max=57 k). \LaTeX templates (.tex/.sty/.cls/.bst) account for 66 % of these files, and 15 % correspond to images or PDFs.

Metadata Level. Metadata either stems from file-level modification timestamps (90 % of submissions contain such data, i.e., were uploaded as an archive preserving this information) or from information embedded in images, PDFs, XML, and other file types that can be extracted using exiftool [69] (93 % of submissions). Out of the timestamps, 98 % are unique, indicating that they are not artifacts of packing submissions for upload to arXiv (per submission, 94 % of files have unique timestamps). Interestingly, 84 % of filtered timestamps coincide with filtered files, confirming a common origin, such as from a ZIP file of a \LaTeX template. For metadata extracted using exiftool, we count a total of 111 M key-value pairs, of which 38 % are unique. Still, non-unique metadata values can carry relevant information, e.g., common software producer names or version numbers.

Content/Text Level. We identify 644 M independent

\LaTeX comments in total; the majority of comments (93 %) are part of required files. Comments are present in 95 % of all considered arXiv submissions. After filtering for unique comments, we obtain 114 M “interesting” comments for analysis that span across 75 % of the considered submissions.

Our analysis not only reveals a large amount of “hidden” information across all three dimensions, but also shows that the majority of this information is sufficiently unique to likely originate from authors directly. Consequently, our analysis in the following focuses on precisely this “interesting” data, as it is substantially more likely to contain sensitive information than repeatedly occurring data (e.g., template files, common producer metadata values, or comments in template files).

6.1.2. Quantifying Frequent Issues. To gain a better understanding of the frequency of specific aspects of “hidden” information, we compare the prevalence and interplay of its three dimensions. As we visualize in Fig. 4, most arXiv submissions contain “interesting” comments (75 %). In contrast, fewer submissions contain unique metadata (56 %) or unique dangling files (25 %). When looking at submissions that contain “hidden” information from more than one dimension, we observe that comments and metadata occur together most frequently (30 %). 18 % of submissions even contain “hidden” information from all three dimensions; most (88 %) contain “hidden” information in at least one of the dimensions.

6.1.3. Differences in Years, Research Fields, and Versions. To better understand potential factors influencing the prevalence of these discoveries, we analyze the prevalence of “hidden” information for (i) changes over time, (ii) differences across disciplines, and (iii) observed variance between different versions of the same submission hosted at arXiv.

Years. As we show in Fig. 4 (left side), the prevalence of “hidden” information in arXiv source files slowly increases over time across the dimensions. This observation particularly holds for the prevalence of dangling files (starting in 2013), and holds for all dimensions since 2017. Perhaps surprisingly, events such as the release or recommendation of sanitization tools or (negative) publicity surrounding “hidden” information have no noticeable impact on the observed trends.

Research Fields. As a multi-disciplinary preprint repository, arXiv is structured into 38 categories. To understand whether the prevalence of “hidden” information shows research domain-specific differences, we aggregate these categories into nine research fields for comparison. As we detail in Fig. 4 (right side), some disciplines, e.g., Math or Physics, feature substantially fewer “interesting” dangling files, potentially because their submissions overall consist of fewer files. Given their high share of submissions (black dots in Fig. 4), they have a substantial impact on the overall results that we report. Thus, covering different research fields when quantifying “hidden” information in arXiv source files is important to gain a comprehensive picture.

In addition to this high-level comparison, we further assess whether computer science researchers, in general, and security researchers (A*/A-/B-/C-ranked conferences), in

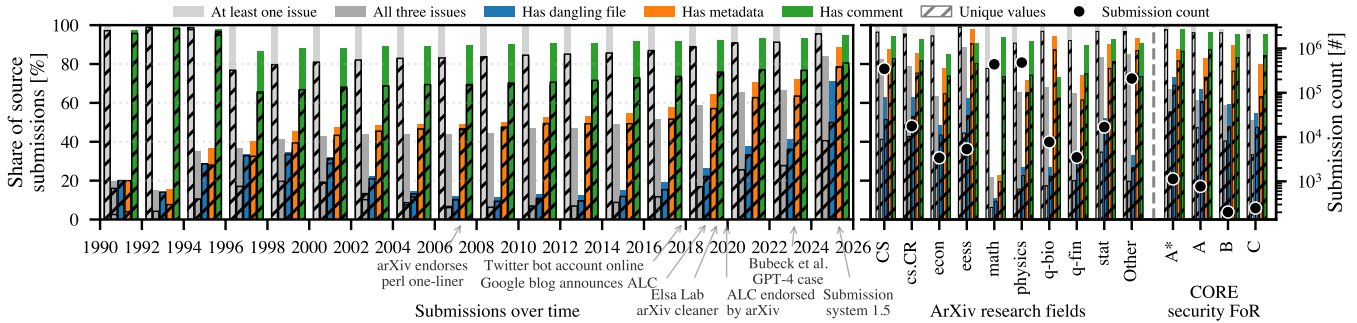


Figure 4. Development of the prevalence of “hidden” information over time shows an upward trend (left) as well as across aggregated arXiv categories and the CORE [83] field of research (FoR) 4604 “cybersecurity and privacy” (right), revealing minor differences across research fields and publication venues.




Figure 5. Subsequent versions rather add unique data over the first version. Only in the long tail, few papers with many versions skew this trend.

particular, exhibit any differences in their behavior (Appendix C details how we match arXiv submissions to published papers by title using dblp [84]). A two-sided Mann-Whitney U test confirms significant differences between all dimensions shown in Fig. 4 when comparing CS (with cs.CR, cryptography and security) against the rest, as well as between cs.CR and the rest of computer science ($p < 0.0001$ in all cases): There is more “hidden” information in computer science papers compared to other disciplines, and even more in security-related papers. Interestingly, the same observation holds when comparing A* security papers with A ones ($p < 0.0001$). Again, higher-ranked venues feature a significantly larger share of “hidden” information. Normalizing issues by the \LaTeX file size (in bytes) to account for length differences does not change this observation. A potential explanation could rather be that computer scientists (and authors of A* papers) have more elaborate source files, while not necessarily paying more attention to “hidden” information. On a positive note, out of 1974 dblp-matched security papers also submitted to arXiv, we discover 243 submissions without any dangling files or other irrelevant content. When also considering the absence of metadata, we identify 70 “perfectly clean” security papers.

Versions. To support the evolution of a preprint over time, arXiv allows superseding a submission with a new *version* (cf. Section 2). While tools for comparing the PDF content of papers over versions exist [85], [86], the impact of new (updated) versions on “hidden” information remains unknown so far. As we illustrate in Fig. 5, many authors update their arXiv-hosted papers over time, even repeatedly, or withdraw them [35], thereby marking them as obsolete (their sources remain publicly available). Intuitively, one reason to submit a new version of a paper could be an attempt to fix or to remove “hidden” information that was accidentally published. Indeed,

individual examples such as a submission initially containing many valid OpenID tokens in dangling files that were later removed with an updated version confirm this suspicion (although this practice is an ineffective mitigation, as the source files of old versions remain accessible). However, at scale, we do not see any consistent indication that authors in general reduce the amount of “hidden” information when updating their preprints with a new version, when comparing the prevalence of “hidden” information between versions in Fig. 5 (in the long tail, the low number of papers with many, ≥ 20 , versions prevents deriving well-founded conclusions).

Takeaway. All three dimensions of “hidden” information are prevalent issues, collectively affecting 88% of arXiv source-code submissions across years, research fields, and submission versions, especially when focusing on “interesting”, i.e., seldomly occurring, content.

6.2. Individually Looking at “Hidden” Secrets

While our initial analysis already establishes the prevalence of “hidden” information in arXiv submissions, so far we only know that the majority (88%) of submissions contain interesting, i.e., unique “hidden” content. We now take a closer look at the individual pieces of “hidden” information.

Guiding Questions. Which individual cases of “hidden” information can we discover? How sensitive are these cases?

To answer these questions, we examine the three dimensions of “hidden” information (file, metadata, and content/text) individually. To this end, we focus on identifying potentially sensitive information disclosure on a *per-submission* basis, i.e., if we identify an issue in multiple versions of the same paper, we only count and report it once.

Where possible, we follow the approach of Leon et al. [87] to label findings as high (H), medium (M), or low (L) impact based on the *potential* severity of them being publicly available. However, this labeling does not imply that all reported findings indeed have the labeled impact.

6.2.1. Category ①: Bulk of Dangling Files. Overall, we discover a total of 12.0 M dangling files. Digging into file and MIME types, we observe the distribution shown in Fig. 6 with a similar distribution across dangling and non-dangling files for frequent file types and a long-tail of rare file types that are predominantly dangling files, including Readme files (3.4%

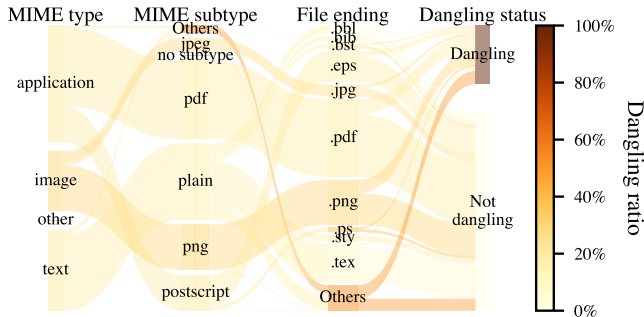


Figure 6. arXiv source files contain a broad range of file types. Some types, e.g., images, are more frequently dangling files than others, e.g., \TeX files.

of submissions, **L**), Python code (0.21 %, **L**), shell scripts (0.21 %, **L**) and other executable files (0.041 %, **L**). We also find research-related data in CSV (0.20 %, **M**) and SQLite form (0.002 %, **M**). 0.007 % of submissions (intentionally) contain “dangling” files within the `anc/` directory [88] (cf. Section 7.3), which we thus do not interpret as dangling.

In general, we cannot reliably assess whether dangling files have been intentionally submitted or not. While 173 k dangling files carry a path such as `supplementary` or `additional_data`, likely indicating intentional submission, 256 k files correspond to directories that contain `backup`, `old`, or similar, instead indicating unintentional submission. Aggravatingly, 40 submissions contain hidden `.nfs` files that are created by NFS filesystems immediately after removal (**M**), and 74 contain complete Git repositories, including editing histories that we consider unintentionally submitted (**M**). Furthermore, we discover configuration files that can carry sensitive information such as API keys or database connection strings in 2774 submissions (**H**, we detail those findings in Section 6.2.3).

Bibliography (`bib`) files are a special case of dangling files, as they are redundant with intermediate `bbl` files that are sufficient for compilation. While `bbl` files comprise only required bibliography entries, we find that 93 % of `bib` files contain entries that are not referenced in the compiled document. Within these dangling `bib` files, we identify two frequent patterns: (a) 8.4 % of submissions reuse bibliographies from prior submissions as is (**L**), and (b) 5.5 % of submissions contain large bibliographies from online repositories, such as ACL Anthology [89], with more than 500 entries (**L**). We find that especially custom-generated and historically-grown bibliographies reveal more information, including note fields (49 % submissions, **M**) and anecdotal citation keys, for example, indicating that authors were scooped by related work (e.g., `oldpaperthatscoopedus`, **M**).

Conference-specific templates are another angle that may reveal information such as submission histories (**L**). Exemplarily analyzing a small, non-representative sample of 115 submissions with dangling `usenix.sty` files where we could match paper titles against dblp records, we find 77 % of these submissions at venues not using USENIX templates (A^* : 61 %, A : 30 %, B : 5.2 %, C : 3.5 %). While we cannot certainly claim that these papers were indeed rejected from USENIX-hosted conferences, these examples suggest that

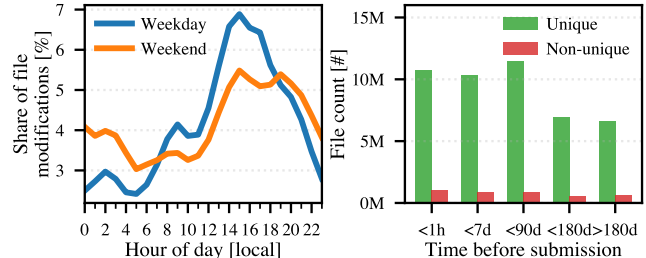


Figure 7. File modification timestamps allow for deriving author behavior from arXiv source files, revealing details of the respective authoring process.

tracing submission histories based on arXiv data could be feasible. We observed and manually validated similar trends for templates of PETS, NeurIPS, CVPR, and ICCV. When identifying venue-specific templates (`.tex/sty/cls/.bst`) as “more than 80 % of template usages occur at the same venue according to dblp”, we count 47 venues and 561 affected submissions that may be traceable this way (**L**). This aspect concerns 0.24 % of the dblp corpus we can study.

Overall, when focusing on files that are unnecessarily submitted to arXiv, we identify both files that in themselves are sensitive, e.g., `.nfs` files or Git histories, as well as files that allow us to infer further, less severe contextual information, such as unused template files hinting at submission histories or `.bib` files revealing collaboration patterns.

6.2.2. Category ②: Insightful Metadata. As authors typically upload a compressed archive that includes original timestamps to arXiv, we can analyze these timestamps and other metadata, which allow for deriving potentially sensitive (**M**) author-specific behavior [75], [90]. 72 % of submissions carry modification timestamps within one hour (35 % within five minutes) before arXiv-recorded submission timestamps. In Fig. 7, we detail the extracted modification times, confirming submission date-based patterns observed in prior work [41] despite extending our analysis to all source files. For creation times, we observe wide ranges, with 83 % of submissions indicating creation dates more than 180 days before submission, reflecting long-term authoring processes. Here, we excluded commonly reused templates, as we only found 38 % of them carrying unique timestamps, whereas most timestamps seem to originate from template providers.

To study file-embedded metadata, we run `exiftool` [69] on all source files. This method yields 62 % of non-dangling files and 61 % of dangling files carrying metadata, most frequently in images (90 % of all images) and PDFs (93 % of all PDFs). While embedded images have been studied before [16], [17], our source-based analysis extends this analysis to also cover metadata of other file types, including dangling files, that have not been part of prior analyses.

Only considering information not accessible from the compiled PDF, we find 11 % of submissions leaking usernames (**M**), 60 % software information (**L**), 2.9 % email addresses (**M**), and 1.4 % specific hardware (**L**), i.e., a significant exposure of sensitive authoring information [75]. In line with related work, we also extract metadata that leaks GPS locations in 7326 of submissions (**H**), 34 % of those findings are limited to dangling files. 2238 sources contain

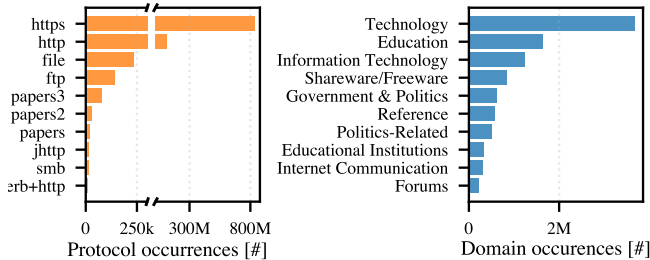


Figure 8. Occurrence of URLs in our dataset grouped by protocol (left) and domain category (right) based on Cloudflare’s domain classification [91].

multiple distinct coordinates; 235 of them have maximum distances of at most 50 km (at least 0.3 km), which we consider commutable, thus possibly enabling the tracking of authors in the temporal and spatial dimensions. Randomly sampling ten of them, we indeed find nine submissions that feature both research buildings and residential areas (H).

6.2.3. Category 3a: Noticeable (Regexable) Content. We now turn our attention to the content/text level within sources that can be efficiently identified through pattern matching. While regex-based methods require further validation and can yield false positives [61], they provide an efficient mechanism for assessing potentially sensitive content at scale. We thus employ an extensive collection of regular expressions [92] to identify concerning content candidates, including cryptographic material, credentials, API keys, network resources, and other patterns related to potentially sensitive information disclosure. For credentials and keys, we employ filtering as recommended in recent related work [93], namely a segmented Shannon-Entropy filter (> 3), besides discarding strictly increasing or decreasing character sequences and a known list of patterns (e.g., aaaa). We run all patterns against all text files in our dataset. Table 2 in Appendix D provides a quantitative overview of our findings. We responsibly disclosed validated critical findings to allow authors to take corrective actions (cf. Appendix A).

Protocol Diversity. We extract and analyze all embedded URLs and IP addresses to understand which external resources authors reference and expose through source files. In total, we identify 127M unique URLs (L) across 905 k submissions and 41 k IP addresses (L) in 19 k submissions. When analyzing the protocol used in URLs, besides the prevalent use of HTTP(S) (98 % of URLs), we also discover references to (legacy) protocols (Fig. 8 left), including FTP servers (29k submissions, 28 % online, 23 % with read access, M), SSH connection strings (349 submissions, L), and other network protocols that may inadvertently reveal infrastructure details or serve sensitive content.

Web References. Leveraging Cloudflare’s domain classification records [91], we further categorize discovered domains in 10/2025 and present their distribution in Fig. 8 (right). Among frequent domains, we find links to cloud storage services, including Google Drive (4272 submissions, M), Dropbox (1542, M), and others, as well as collaborative platforms such as Overleaf (26k, M) and Google Docs (3948, M). Manually accessing a randomly sampled subset of these links reveals that some of the linked services are

publicly accessible, i.e., they might serve sensitive content unintentionally. Overall, 4283 embedded links were—at the time of the conducted classification—labeled as “file sharing” by Cloudflare. Moreover, as classified in 10/2025, we discover 587 links related to questionable content (incl. CIPA filter, profanity, and pornography, L); 7 links serve violent content (according to Cloudflare, subject to change).

Content Classification. Exemplarily looking at the content behind found web references (cf. Appendix A for ethical considerations), we observe that 1119 Google Docs links grant viewing (M) or 699 editing (H) access, with at least 200 cases exposing sensitive content (according to manual analysis, H). This content covers commonly-confidential documents from the authoring process (reviews, cover letters, rebuttals, and revision logs), todo lists with both ticked and unticked items, meeting minutes (in one case, including links to Zoom recordings), student assignments, and even a shift schedule, as well as a stock and order list. For other content types, e.g., datasets, evaluation results, figures, or plots, we cannot assess the sensitivity without additional context. We also find several Google Docs links that allow us to transitively reach other web resources not referenced in the source files, resulting in broader information disclosure. Alarming, we also came across 18 cases where we obtained access to survey data of study participants (H).

Security-related. Further analyzing irrelevant content for clearly security-critical content, we discover 265 API tokens in 128 submissions (26 % dangling, H), 4 private keys in 4 submissions (50 % exclusively in dangling files, H), and 171 generic passwords in 82 submissions (25 % dangling, H) after manual validation. Their placement within the source files indicates that the authors likely did not intend to distribute them publicly. These discoveries align with related work on secret detection in software repositories [57], [58], [60] and highlight that the problem extends beyond traditional code-sharing platforms to academic preprint services.

Malware Detection. Using YARA’s established pattern-based approach for detecting malware [94], we attempted to discover malicious (binary) patterns within the source files. However, upon manual validation, all hits turned out to be false positives. Thus, this straightforward approach for identifying malicious content did not reveal hazardous risks.

“Hidden” LLM Instructions. We also investigated the source files for indicators of problematic authoring activities. Recent work has highlighted concerns about LLM-generated content being used to influence peer review [95]–[97], prompting us to search for patterns that resemble prompts or instructions to language models. Apart from confirming prior findings of instructions within the PDF [97], we specifically looked for prompt injection patterns in commented-out sections, i.e., instructions that may have deliberately been part of the paper before, e.g., during peer review. In fact, we manually confirmed nine submissions containing such instructions. Besides, 537 submissions embed hidden text fragments that match common LLM response disclaimers like “As an AI language model” in commented-out sections, which indicates generated content [98] (L). However, their actual purpose and impact on the document’s text remain

unclear without deeper semantic analysis.

Other Noticeable Content. Similarly, searching for review-indicating keywords such as “reviewer,” “rebuttal,” or phrases such as “wrong” and “not correct” yields 3.2K submissions containing such terms in commented parts, potentially indicating responses to peer feedback that authors forgot to remove before submission (L). We also discover that 269 submissions actively use L^AT_EX packages designed for content redaction, such as `cancel` and `pdfprivacy`, suggesting attempted but ineffective sanitization (H). Further, within dangling files, we encounter LICENSE files in 5602 submissions (L), which may be incompatible with the submission’s license, potentially unintentionally open-sourcing differently-licensed material. This finding raises questions on intellectual property considerations and whether authors understand the implications of distributing their research under specific open-source licenses through arXiv.

Finally, applying profanity detection patterns [99], we discover offensive language in citation keys, e.g., `fuck1` to `fuck7`, appearing in 101 submissions (L). This finding carries implications also for generated PDF documents, as citation key names and reference labels become accessible in PDF metadata when using the `hyperref` package—which arXiv enabled by default in its original submission system (1.0) [100]. In this context, we identify 42M custom—likely irrelevant—bibliography entry fields that deviate from standard BibT_EX schemas, mostly containing editorial comments or annotations that were likely not intended for public distribution: 29M entries across 150k submissions reveal author comments, local file paths, and behavior, i.e., creation and modification dates as well as ratings of papers (L).

6.2.4. Category 3b: Author Comment Analysis. Moving beyond noticeable content that can be identified via pattern-matching, we now strive to identify more subtle aspects of “interesting” content, i.e., unique author comments.

Methodology. To better understand these individual comments at scale (see Section 4 for what we consider a comment), we task an LLM with classifying the content, building on promising results for text classification and sentiment analysis across diverse domains [101]–[103]. Given the potentially sensitive nature of the content, we utilize a locally-hosted Qwen2.5–72B multi-language model to avoid transmitting any content to third-party services (cf. Appendix B for a detailed discussion of our LLM usage considerations). According to `lingua-py` [104], 92% of comments are in English language, followed by Latin (2.9%, often misclassified math commands), European (2.5%), and Asian (0.6%) languages. Overall, >96% of comments correspond to languages supported by the model [105]. We classify comments into six categories using in-context learning with at least two examples per category: *formatting* (navigation aids and visual separators), *academic text* (commented-out paragraphs or alternative phrasings), *conversational* (rather informal notes between coauthors), *L^AT_EX markup* (table structures or mathematical expressions), *todo* (task lists and open items), and *bibliography* (citation-related notes). Additionally, we instruct the model to flag comments

containing potentially private or sensitive information. On a manually labeled test dataset of 500 random comments (stratified by LLM-predicted class), this method achieves a macro-averaged F1-score of 0.83 for assignment to categories, and 0.71 for flagging sensitive content.

LLM-based Classification. Our evaluation shows that L^AT_EX markup comments are most prevalent (30%, 92% of submissions, L), anecdotally containing table formatting, mathematical expressions, or disabled code. The second most common category is *academic text* (44%, 76% of submissions, L), comprising commented-out paragraphs, alternative text formulations, or previous versions of arguments. Manual inspection of selected cases reveals anecdotal evidence of conflicting lines of argumentation between published and commented text, though we leave a systematic analysis of semantic contradictions for future work. *Conversational* comments rank third (7.9%, 72% of submissions, M) and frequently contain discussions among coauthors, including critical judgments about text quality, suggestions for improvement, or strategic considerations about presentation—content presumably not intended to be public. Less common categories are *formatting aids* (1.7%, 64% of submissions, L), *bibliography-related comments* (9.1%, 62% of submissions, L), and *todo items* (0.98%, 44% of submissions, M). For the latter category, we cannot automatically assess whether these tasks were ultimately addressed or not. However, manual examination reveals cases where todos explicitly mention methodological limitations or acknowledge weaknesses that are not disclosed in the published text.

The model also flags 7.0% of comments as sensitive, including PII, opinions about reviewers/related work, author discussions, or credentials (M–H). Via bootstrapping ($n = 10000$, 95% confidence) of the labeled dataset, we estimate that at least 3.6% indeed contain sensitive information.

Comparing the different comment mechanisms (cf. Section 4), we find that most comments are standard line comments using % (70%). We also identify comment blocks within comment environments (15%), content in skipped branches of conditional statements (0.14%), arguments discarded by custom commands (8.8%), post-document content (0.085%), and text within special environments (0.055%). Cross-referencing these mechanisms with our semantic categories reveals some notable patterns: *conversational* comments appear disproportionately often in comment environments (15%), whereas 83% of *todo* comments are simple line comments, suggesting that this mechanism is primarily used for tracking tasks during collaborative writing.

Disclaimer. These numbers should be taken with a grain of salt as (a) LLM classification can introduce misclassification, partly as we use a weaker self-hosted model for security reasons (cf. Appendix B) and (b) the perception of what is sensitive varies substantially across individuals [87]. Still, we believe that this approach is the best one available as manual annotation not only results in annotator bias but also fails to scale to millions of comments.

Takeaway. Striking individual cases span a wide spectrum, including cryptographic material, credentials and PII of survey participants, editable online documents, (GPS

TABLE 1. COMPARISON OF AVAILABLE TOOLS FOR SOURCE FILE CLEANING (SORTED BY TECHNOLOGY AND ENDORSEMENT BY ARXIV [6]). NONE OF THESE TOOLS IS ABLE TO RELIABLY SANITIZE ALL TEST CASES. REGARDLESS, A SUBSET OF AUTHORS SEEMS TO APPLY THEM PRIOR TO SUBMISSION.

Name	Technology	Claimed Features	Last Update	Comment Cleanup Tests								Sanitization		Comment Cleanup Tests:		
				①	②	③	④	⑤	⑥	⑦	⑧	⑨	“Beneficial”	Breaks [§]	① Inline comment removal	
perl one-liner [6]	regex-based	comments	11/2005	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	95 %	4.2 %	② \%
arxiv_latex_cleaner (ALC) [21]	regex-based	• dangling files • comments	03/2026	✓	✗	✓	✗	✓	✓	✗	✗	✗	✗	91 %	19 %	③ Comment environment
latexindent.pl [64]	regex-based	comments	02/2026	✓	✗	✗	✗	✗	✗	✗	✗	✗	98 %	28 %	④ Retain comments within special environments	
arXiv Cleaner [65]	• pdflatex recorder • latexpand [106]	• dangling files • comments	09/2019	✓	✓	✗	✗	✗	✗	✗	✗	✗	3.7 %	9.2 %	⑤ Out of document removal	
Sub. Sanitizer & Flattener [66]	snapshot [107]	dangling files	06/2022	✗	✗	✗	✗	✗	✗	✗	✗	✗	80 %	80 %	⑥ \iffalse/if0 handling	
pandoc [108]	tokenization	n/a	03/2026	✓	✓	✓	✓	✓	✓	✓	✓	✓	69 %	85 %	⑦ Custom \if removal	
ALC-NG [26] (<i>this work</i>)	• pdflatex recorder • exiftool [69] • tree-sitter-based [80]	• dangling files • metadata • comments	04/2026	✓	✓	✓	✓	✓	✓	✓	✓	✓	98 %	15 %	⑧ No arg. command cleaning	

[§] (i) tool crashes, (ii) paper cannot compile, or (iii) mismatch in pixel-perfect comparison
 ✓ cleans \LaTeX source files successfully, ✗ unsuccessful (attempt), ⚠ requires manual action for sanitization, and † tool crashes, i.e., failed test

metadata from authors, unfinished tasks on todo lists, traceable submission histories, and collaboration artifacts. Failed-to-delete NFS files and evidence of ineffective censorship emphasize the authors’ intent to hide such information.

7. Mitigation Through \LaTeX Sanitization Tools

At first glance, the prevalence of “hidden” information might appear surprising, as different tools to sanitize (or “clean”) \LaTeX sources—before submitting them to arXiv or similar services—and thus mitigate the identified issues are readily available, some even actively advertised by arXiv [6].

To establish a structured understanding of why these existing \LaTeX sanitization tools fail to address the issue of “hidden” information in arXiv submissions, we first survey the landscape of cleaner tools (Section 7.1) and then evaluate their performance in cleaning \LaTeX sources as well as estimate their usage frequency (Section 7.2). To address the inherent shortcomings of existing sanitization approaches our analysis identifies and to provide mitigations for the problem of “hidden” information uncovered by our work, we present ALC-NG, our new sanitization tool (Section 7.3).

7.1. Landscape of \LaTeX Sanitization Tools

Different so-called *sanitization tools* promise to “clean” \LaTeX source files to prevent issues stemming from “hidden” information when preparing submissions, e.g., to arXiv. In the following, we provide a structured and comprehensive overview of the various sanitization approaches and provide a first assessment of their cleaning capabilities.

To identify tools and software designed for cleaning arXiv source files, we systematically searched for projects on GitHub using the keywords `arXiv cleaner` OR `arXiv sanitizer`. We further considered a compilation of different approaches on StackExchange from 2012 [109]. We then selected representative approaches, one per underlying technology, based on popularity (stars on GitHub) and compiled a detailed comparison of these tools in Table 1.

Out of these sanitization tools, the following warrant further discussion due to their popularity or unique approach. *arxiv_latex_cleaner* (ALC) [21]. Arguably, this cleaner developed by Google is the most prominent one, operating

both on the *file level* and the *content/text level*, as evidenced by various recommendations [110] and arXiv’s endorsement since 12/2019 [6]. Different academic works even apply it to preprocess input data as part of their research methodology in the context of LLM/AI research [111]–[113]. Despite ongoing efforts to ensure correct behavior [114], several open issues related to the accurate detection (and removal) of comments, i.e., irrelevant content, remain [115], [116].

Submission Sanitizer & Flattener [66]. In contrast, the Submission Sanitizer & Flattener focuses specifically on the *file level*, i.e., solely removing files that are not needed, but not considering the irrelevant content and text contained in the remaining files. As such, it differs quite substantially from the other tools in our comparison. To identify which files are needed to build a \LaTeX document, an additional build process is required using the `snapshot` package [107].

pandoc [108]. While not advertised as a sanitization tool, `pandoc`, as a universal document converter, can also clean \LaTeX sources at the *content/text level*. However, it massively reformats the sources, adds new \LaTeX instructions, and breaks compatibility with common templates (e.g., `acmart`), severely challenging its usage, let alone its automated usage.

Given this diverse landscape of sanitization tools, we sought to better understand their capabilities. To this end, using synthetic experiments, we assess their ability to remove irrelevant content alongside the four concepts for embedding comments in \LaTeX documents we derived in Section 4 (A) comments, (B) outside document, (C) non-taken `if` branches, (D) unused arguments of commands). We created eight distinct test cases (①–⑧) that comprehensively cover these four concepts (cf. Table 1), including straightforward and more subtle ways of irrelevant content as well as content that might be falsely detected as a comment (e.g., text within `verbatim` environments). Additionally, we test whether sanitization tools replace calls to `BIBTEX` with the content of the corresponding `bbl` file (⑨). Our research artifact [25] contains a minimal working example that covers these cases.

The results of our synthetic comment cleanup tests listed in Table 1 show that none of the sanitization tools perform satisfactorily in terms of cleaning coverage, despite arXiv endorsing two of them [6]. Even among tools that utilize the same technology, e.g., regex-based ones, we observe notable differences in their cleanup coverage. Surprisingly,

the arXiv Cleaner [65] even crashes when encountering a basic `comment` environment in the source file. Additionally, to remove certain irrelevant content, some sanitization tools require manual interaction, which impairs usability for inexperienced \LaTeX users. Moreover, all cleaners but ALC [21] require the user to explicitly identify the primary \LaTeX file, potentially constraining the ease of use. This situation is somewhat unexpected since arXiv’s submission system can process submissions automatically, even if it means falling back to the arXiv-specific `00README(.XXX)` file [78].

All available cleaners exclusively cover the *file level* and (certain aspects of) the *content/text level*, ignoring (sensitive) content at the *metadata level*. Hence, they entirely disregard a complete dimension (cf. Section 4) to the problem of “hidden” information in \LaTeX sources.

7.2. Quantifying Sanitization Performance

While our analysis of \LaTeX sanitization tools (cf. Section 7.1) provides a *theoretical* understanding of their capabilities (and limitations), a *practical* quantification of their sanitization performance on actual arXiv submissions is still missing. Thus, we now evaluate (a) whether the application of sanitization tools would have removed additional “hidden” information in arXiv submissions, (b) which “hidden” information still remains after sanitization, (c) whether sanitization tools have been applied on sources prior to submission, and (d) how large the potential for storage savings at arXiv is if sanitization tools were applied consistently.

Methodology. To answer these questions, we run the identified cleaners (Table 1) on our dataset (Section 5). To establish a reliable baseline for dangling files, we monitor file system accesses during the build process and classify files that have not been accessed as dangling. For the build process, we utilize the most recent \TeX Live version in Ubuntu 24.04 (2023) and `latexmk` [117] (v4.83). Due to breaking changes over time [8], both in \TeX Live and author-included \LaTeX packages, we are unable to successfully compile *all* \LaTeX sources. Still, we can compile—and thus evaluate the different sanitization tools on—2.32 M papers (85 %).

Sanitization Performance. To quantify whether the application of sanitization tools would have been beneficial, i.e., removes additional “hidden” information, we apply the different cleaners and evaluate whether they reduce the amount of dangling files and comments for a given paper. Table 1 reports (a) the share of papers for which a sanitization tool alters the sources, with the paper still compiling after application (“*beneficial*”—theoretical benefit, albeit no guarantee of full sanitization coverage or accurately retaining visible content), and (b) the share of papers where the tool *breaks* the sanitization attempt, either via tool failure, removal of essential content, or other issues that visually alter the document. Exemplarily focusing on ALC as the most prominent tool, 91 % of the evaluated submissions (2.12 M) would “benefit” in terms of (additionally) removed “hidden” information. In detail: For 50 %, the number of dangling files decreases; similarly, for 91 %, the number of comments decreases. ALC supports manually specifying commands

that should be removed. Even though we expect laypersons to rarely make use of this manual feature, we report numbers for both behaviors in Table 1 as well. Importantly, these results do not entail that ALC perfectly sanitizes those papers. Rather, it merely *reduces* the amount of “hidden” information. Sometimes, it also *overremoves relevant* content, leading to an inflated sense of beneficialness, i.e., “breaking” the content. Similar observations hold for the other sanitization tools.

Especially when identifying (and removing) dangling files, relying on regexes (as done by ALC) comes with the inherent risks of a heuristic approach, i.e., potentially leading to over- or underremoval. As a result, \LaTeX sources do not compile anymore, or dangling files—other than intended—are not removed. Quantifying this issue, we also visualize the difference between using the file system access-based approach (cf. methodology) as a reliable baseline and ALC in Fig. 9 (Appendix D). Similar issues arise for regex-based comment removal, where special environments, such as `listing` or `verbatim`, are frequently mismanaged, leading to “clean” sources with altered content.

Consequently, while existing tools are beneficial in that they remove sensitive information, they also come with the risk of breaking \LaTeX sources and altering content, albeit not comprehensively removing sensitive information.

Estimating Attempted Sanitization. To roughly quantify this false sense of security, we check whether applying a sanitization tool’s most recent version results in unchanged sources, which may indicate previous use. We find 106 k submissions without any comments and/or dangling files, which authors might have sanitized (manually or using a sanitization tool). Still, the large number of papers where cleaners are considered “beneficial” highlights that only few authors were able to successfully apply them. Trivial line comments, which every sanitizer passing ❶ removes accurately (cf. Table 1), in 94 % of evaluated submissions suggest that no comprehensive attempt was made. Considering the failure to handle ❷ as a fingerprint of ineffective ALC usage, we find 538 uses (0.10 % of submissions since its release in 2018 [118]), i.e., a minuscule prevalence of confirmable sanitization with an arXiv-endorsed tool [6].

Remaining Sensitive Information. We now focus on any remaining sensitive information *after* applying sanitization tools. In addition to the underremoval of dangling files, we identify thousands of remaining, i.e., underremoved, *comments* after running ALC. Within those comments, our regex-based detection (cf. Section 6.2.3) reveals 362 matches, e.g., including 7 (sensitive) links to Overleaf, highlighting the need for “perfect” removal. Consequently, current sanitization tools might create a false sense of security by potentially *not* removing all occurrences of sensitive information.

Storage Overhead. Besides unintentionally disclosing “hidden” information, missing out on potential for sanitizing unnecessary information from arXiv submissions also puts a *permanent* storage burden on arXiv. For the set of 2.32 M submissions we can compile, we identify 658 GB (22.9 %) that could be freed by accurate sanitization of the analyzed version only; 97 % account for dangling files, 0.4 % for metadata, and 2.3 % for content in required files.

7.3. ALC-NG – A Reliable L^AT_EX Sanitization Tool

Motivated by the shortcomings of current L^AT_EX sanitization tools and following our ethical obligation to provide a reliable mitigation for the threat posed by “hidden” information in L^AT_EX sources, we open-source our “Next-Generation” sanitization tool, called ALC-NG [26] (See: <http://alc-ng.de>).

With this tool, we combine established best practices for each dimension of “hidden” information (cf. Section 4), i.e., pdflatex recorder for dangling file detection, exiftool for metadata removal, and a tree-sitter-based detection of irrelevant content (comments). While arXiv Cleaner [65] also applies pdflatex recorder, state-of-the-art sanitization tools do not yet use exiftool or a tree-sitter grammar (cf. Table 1).

① *File Level*. To reliably identify files that are required to compile a L^AT_EX document, we utilize *pdflatex recorder* (similar to arXiv Cleaner [65]). Using the access timestamp (*atime*) of the file’s last read [119], we validated the correctness of this approach on the 2.32M compilable papers by comparing its output with read accesses on disk and did not discover any discrepancies. Furthermore, we use the 00README file [78] as well as the directory `/anc/` for ancillary files [88] to identify files *not* to remove.

② *Metadata Level*. To sanitize metadata, we utilize *exiftool* [69] since it is known for reliably removing metadata across various file types, including PDFs and images. Optionally, we support suppressing file modification times to hide corresponding sensitive behavioral metadata further.

③ *Content/Text Level*. To reliably identify and remove unnecessary content, we re-purpose our abstract syntax tree-based detection logic (cf. Section 4). As we summarize in Table 1, this approach comprehensively covers all aspects of our synthetic comment cleanup tests.

To validate the correctness of ALC-NG’s sanitization, i.e., ensure that the removal of files, metadata, and content does not change the resulting compiled PDF, we compare the original and the sanitized PDFs using *pdfium* [120] to ensure *pixel-perfect* versions, i.e., visually-identical PDFs. We get visually-identical PDFs for 87% of the compilable submissions *after sanitization*; as opposed to 96% for the perl one-liner as the best cleaner for this metric, but subject to underremoval (little sanitization coverage), and 88% for ALC, albeit for a lower number of submissions that remain compilable *after sanitization*. Applying the current prototype of ALC-NG on the 2.32M compilable submissions shows that 98% (2.27M) benefit from its use by excluding dangling files and/or cleaning irrelevant content: ALC-NG comprehensively removes the different dimensions of “hidden” information within L^AT_EX sources (cf. Section 4). Despite its proof-of-concept implementation, it outperforms the surveyed sanitization tools in terms of coverage (cf. Table 1), while promising great ease of use even for inexperienced L^AT_EX users, laying the foundation for broad adoption.

Consequently, with ALC-NG, we empower authors to reliably sanitize their L^AT_EX sources before submitting them to arXiv (or other third parties) and thus mitigate the various threats posed by “hidden” information identified in this paper.

8. Discussion and Implications

After concluding our analysis of “hidden” information and sanitization tools, we now focus on permeating our findings. Specifically, we discuss key implications (Section 8.1), perception in the community based on a survey among affected arXiv authors (Section 8.2), limitations of our work (Section 8.3), arXiv-specific aspects (Section 8.4), and recommendations for stakeholders (Section 8.5).

8.1. Putting Our Findings into Perspective

Throughout our work, we considered the guiding questions: *How prevalent are the three dimensions of “hidden” information in source files of preprints? Which individual cases of “hidden” information can we discover? How sensitive are these cases?* We now revisit these questions.

General Assessment. Our analysis uncovered a wide range of “hidden” information. We discovered information with immediate confidentiality and security implications, as well as details that can be used to deduce behavioral aspects, e.g., timestamps or sanitization attempts. Hence, the issues we identify vary greatly in severity. Even seemingly trivial attempts to mitigate the issue—accurately removing dangling files prior to submission—would have removed 20% of links, 33% of GPS locations, and 5.9% of unique comments. On another note, perhaps counterintuitively, papers appearing on both arXiv and in A* - and A-ranked security conferences disclose, on average, more “interesting” and sensitive information than their counterparts from other areas, despite security researchers’ presumed awareness of such risks.

In Appendix D, we further complement this paper by discussing “hidden” information in our *own* arXiv papers. Moreover, we revisit the GPT-4 paper (cf. Section 1) to assess whether that case could have been easily prevented.

Risk of Breaking Blindness in Peer Review Processes. Our findings have implications beyond preprint servers, potentially affecting (double-blind) review and publication processes where authors supply source files. Particularly, beyond the well-known implications of metadata, reviewers or infrastructure providers could exploit additional “hidden” information once more reviewing and publication processes require the submission of L^AT_EX source files. Relatedly, the rise of collaborative L^AT_EX editors [121] will further increase the usage and distribution of L^AT_EX sources.

8.2. Responsible Disclosure and Author Survey

Throughout our research, we responsibly disclosed validated security-critical findings to 2660 authors of affected papers, informing them about our findings and allowing them to take corrective actions (cf. Appendix A for more details on our considerations surrounding the responsible disclosure).

We complemented our disclosure notifications with an overview of other sensitive findings (e.g., potentially sensitive L^AT_EX comments) and a link to a survey to assess awareness and practices regarding L^AT_EX sources when submitting to arXiv. We received 112 complete survey responses (response

rate: 5.3%). The largest group of respondents were assistant (29%) or full (16%) professors, followed by permanent research staff (28%) and postdoc/PhD students (18%). 49% of respondents claimed more than ten arXiv submissions, and only 41% stated awareness of source file availability and its implications before receiving our disclosure notice.

Sensitivity. 43% of respondents agreed that the disclosed information is sensitive. However, when asking specifically about the dimensions of information leakage, only 21% confirmed that no sensitive information was contained in their sources, indicating disagreement and uncertainty about what constitutes sensitive information (cf. disclaimer in Section 6.2.4). For the other respondents, *conversational* comments were considered sensitive by 39%, and *todo* comments by 38% of respondents (cf. Section 6.2.4). For the 21% reporting no sensitive information, we find significantly fewer dangling files (-74%), metadata values (-40%), and unique comments (-36%, Mann-Whitney-U test: $p \leq 0.03$ in all cases). This finding suggests that authors possess some intuition regarding sensitive hidden information. 9 email respondents stated that disclosed online documents did not contain sensitive information, but 7 still restricted access. By March 2026, we discovered 18 updated submissions that followed our disclosure, indicating cleanup attempts. Irrespectively, the affected, old source files remain accessible.

Perceived Helpfulness. 57% of respondents indicated that the disclosure was very helpful (79% when including participants who found it rather helpful). Separate from survey participation, 51 authors replied positively to our disclosure via email. In contrast, two respondents reported false positives from our detection, requesting that we should have invested manual effort to verify the issue—which we thoroughly did (cf. Section 6.2.3).

Experience with Cleaners. Among respondents, 28% had heard of cleaning tools listed in Table 1, 12% had tried them, and 11% report using at least one cleaner regularly. Regarding cleaning practices, 8.9% reported applying (manual) cleaning before the submission, which we quantify as reducing 36% of unique comments in the notified-upon submissions; however, *all* of these submissions still had issues. Seven authors responded that, following our disclosure, they uploaded a new, cleaned version—*not* realizing that old versions and sources remain accessible.

Potential Biases. We acknowledge potential bias in our survey responses, since we primed participants immediately before participation by indicating that we considered their sources to leak potentially sensitive information. Moreover, we only contacted authors who disclose publicly-accessible Google Docs URLs or credentials (manual curation) within irrelevant content, disregarding (more careful) authors without such findings. Likewise, authors who did not consider our findings sensitive might have been less inclined to respond to our survey and/or did not expect other sensitive findings. Still, we received text responses containing more details than disclosed in our email, indicating that at least some authors carefully reflected on our disclosure campaign.

8.3. Limitations and Implications

While our methodology enables the study of “hidden” information in arXiv sources, it is not without limitations.

These limitations mainly stem from the scale of the research problem, which demands different assessment directions as evident in the diversity of our analysis (cf. Section 6.2). First, we pursue a best-effort approach when analyzing the sources and compiling paper PDFs, which comes with some minor limitations: (i) 1 source file archive is corrupt. (ii) We limit both parsing and compiling of source files to 5 min and kill the compiler once its logfile exceeds 2 GB, resulting in 388 prematurely terminated analyses. (iii) We ignore parsing errors when processing the sources using the tree-sitter grammar (cf. Section 4), affecting 681 k submissions, resulting in, on average, 9829 bytes of disregarded content that may contain “hidden” information.

Second, our comment detection within \LaTeX files revealed 644M occurrences. While ideally we would have taken a manual look at *all* individual comments, this approach is clearly infeasible. Instead, we combined our definition of “interesting” content, various regex detection methods, and an LLM-based analysis of author comments to gain a comprehensive picture. Arguably, this reasonable best-effort approach may have missed important sensitive details and other issues. As such, the details presented in this paper “only” provide an accurate lower bound concerning the prevalence of “hidden” information, potentially missing more findings.

For the implications of applying an LLM-based content classification, we refer to our disclaimer in Section 6.2.4 and our LLM usage considerations in Appendix B.

Likewise, as our assessment could only focus on source files that remain publicly accessible, we could not consider submissions and versions where sources have purposefully been removed. Thus, our version-related conclusion (cf. Section 6.1.3) may be skewed, since very sensitive source files may have actively been removed. Only if arXiv persisted those sources, they could still conduct such an analysis internally. Irrespectively, third-party researchers cannot address this data collection-inherent limitation for very good reasons.

Finally, for our evaluation of existing sanitization tools (cf. Section 7.1), we constrained our analysis to arXiv submissions that we could compile in a modern \LaTeX build environment (cf. Section 7.2). As a result, our reported numbers cover 85% of all submissions with sources. For the remainder, evolving \LaTeX compilers and breaking changes in packages prevent compilation [8]. Still, by covering the majority of submissions with sources in our analysis, we are confident that it accurately reflects the performance of the different sanitization tools, especially considering that these tools will (only) be applied to new arXiv submissions and thus will compile within a modern \LaTeX build environment.

8.4. Identified arXiv-specific Peculiarities

Given our focus on arXiv, we also noted and considered some of its specific peculiarities, as detailed below.

Implications of the new Submission System. In April 2025, arXiv introduced a new submission system (1.5). We can easily track the usage of this submission system through the values of the `Producer` metadata key. Even though it introduces notable changes to submission processing, we do not observe any immediate implications on our findings: E.g., the new system now specifically prompts submitters to remove dangling files on the root-level, while other directories and their dangling files remain unaffected. Fig. 10 (Appendix D) shows a zoomed-in version of our analysis of “hidden” information for 2025, indicating no substantial changes through the switch to the new submission system.

arXiv-specific Cleaner Functionality. ALC-NG—different from previous sanitization tools “for” arXiv—supports arXiv-unique functionality. Thus, we support both `00README(.XXX)` files [78] (used to override defaults) and `anc/` directories [88] (to intentionally add ancillary files).

Identifying Exploitation of Findings. Complementing our assessment of arXiv sources, we were interested in identifying whether other parties investigate such “hidden” information as well. To this end, we added “honey pot” URLs to an otherwise benign arXiv submission within the CS category. More specifically, we added multiple URLs to the paper across all three dimensions of “hidden” information (cf. Section 4) and monitored HTTP(S) requests. However, since its submission to arXiv in 05/2025, we have only recorded accesses to URLs that are embedded within the compiled PDF file. URLs that are solely part of the sources have not been accessed yet. While limited in scope, this observation suggests that, since deployment, no large-scale analysis with a similar methodology as ours (involving accessing discovered URLs) processed our “honey pot” submission in detail.

Meta Aspects. Given that the majority of arXiv submissions features “hidden” information, partially leaking very sensitive details, future work should focus on better educating and thus supporting authors. This aspect is particularly important since arXiv is currently not considering additional changes related to the (retrospective) sanitization of “hidden” information. Once a tool like ALC-NG has matured sufficiently, arXiv might consider integration into its submission system. Beyond the reliable removal of “hidden” information, a substantial reduction of storage needs (cf. Section 7.2) could be another driver for such a change.

8.5. Recommendations for Stakeholders

Based on our findings and considering the prevalence and severity of “hidden” information in source files of preprints, we derive recommendations for affected stakeholders.

Authors and Institutions. To ensure confidentiality of sensitive information and intellectual property, authors must ensure that they remove all comments and dangling files from their source files prior to submission to a preprint repository, preferably using a thorough sanitization tool such as ALC-NG [26]. Although not in the main focus of our work, this sanitization should include removing sensitive metadata from included PDF files and images (cf. Section 6.2.2), using either dedicated tools [69]–[71] or (manually) orchestrated

through ALC-NG. From an institutional perspective, to protect affiliated authors and intellectual properties, institutions should offer guidelines and establish precise policies on the release of \LaTeX sources. Already affected authors and institutions, i.e., those ones that (unknowingly) released not (properly) sanitized source files in the past, are likely unable to fully remove unintentionally-disclosed information (cf. Appendix A). Consequently, corrective actions should concentrate on revoking permissions for, e.g., exposed API keys and enforcing access control for affected URLs.

Preprint Repositories. Complementing efforts by authors and institutions, preprint repositories should raise awareness for the risks resulting from “hidden” information and enforce sensible behavior. For example, by incorporating checks on irrelevant content that sanitization tools would remove, preprint repositories could issue warnings for detected potentially sensitive “hidden” information in uploaded source files and automatically discard dangling files. For already published source files, preprint repositories could proactively check for indicators of sensitive “hidden” information and work with authors to expurgate affected submissions.

To ease the understanding, we provide an executive summary of our findings at <https://arxiv.comsys.rwth-aachen.de>.

9. Conclusion

In this work, we systematically assessed the issue of unintentionally-disclosed information in source files of 2.7 M arXiv submissions. Our findings reveal “hidden” information in almost every submission, albeit with greatly varying sensitivity. Apart from unused template files that put unnecessary storage burden on arXiv, we further discover scripts, research data, and even entire Git repositories. Additionally, comments in \LaTeX sources reveal, e.g., author conversations or todo items—for some of those comments, we are certain that the authors did not intend to disclose them publicly. Alarmingly, our findings also include URLs without any access restrictions to other resources (e.g., Google Docs), security tokens, and private keys. Ultimately, the issue concerns three dimensions: dangling files, embedded metadata, and irrelevant content, e.g., in comments.

Considering possible mitigations, our work reveals that current \LaTeX sanitization tools leave room for improving coverage and accuracy (resulting in both under- and over-removal). Combining technologies known to perform well with a previously unconsidered approach to detect \LaTeX comments, we develop ALC-NG, a reliable cleaner for source files. Our findings further show that better education among researchers is needed to holistically address the issue of unintentionally-disclosed information within source files.

Acknowledgment

Thank you to arXiv for use of its open access interoperability and for supporting our research activities. The authors thank Cloudflare (domain classification) and NVIDIA (compute activities) for supporting this work, and Muzamil

Bashir for his initial, prototypical validation of this research. We also thank the anonymous reviewers and our shepherd for their valuable feedback and suggestions.

References

- [1] G. McKiernan, “arXiv.org: the Los Alamos National Laboratory e-print server Available to Purchase,” *IJGL*, vol. 1, no. 3, 2000.
- [2] citeordie, “Let’s discuss: is arXiv always good?” <https://archive.is/ZQ5u2>, 2015.
- [3] Cornell Tech, “arXiv.org e-Print archive,” <https://arxiv.org/>, 1991.
- [4] B. Xie *et al.*, “Is preprint the future of science? A thirty year journey of online preprint services,” arXiv:2102.09066, 2021.
- [5] arXiv, “Monthly Submissions,” https://arxiv.org/stats/monthly_submissions, 2014.
- [6] arXiv, “Requiring TeX when possible,” <https://arxiv.org/help/faq/whytex.html>, 2005.
- [7] J. Gehrke *et al.*, “Overview of the 2003 KDD Cup,” *SIGKDD Explor.*, vol. 5, no. 2, 2003.
- [8] J. Tan and M. Rigger, “Inconsistencies in TeX-Produced Documents,” in *ISSTA*, 2024.
- [9] Y. Deng *et al.*, “Challenges in End-to-End Neural Scientific Table Recognition,” in *ICDAR*, 2019.
- [10] N. Jiang *et al.*, “LATTE: Improving Latex Recognition for Tables and Formulae with Iterative Refinement,” *AAAI*, vol. 39, no. 4, 2025.
- [11] S. Brinn *et al.*, “A framework for improving the accessibility of research papers on arXiv.org,” arXiv:2212.07286, 2022.
- [12] S. Byers, “Information Leakage Caused by Hidden Data in Published Documents,” *SP*, vol. 2, no. 2, 2004.
- [13] T. Aura *et al.*, “Scanning electronic documents for personally identifiable information,” in *WPES*, 2006.
- [14] A. Castiglione *et al.*, “Taking advantages of a disadvantage: Digital forensics and steganography using document metadata,” *JSS*, vol. 80, no. 5, 2007.
- [15] C. Alonso *et al.*, “Disclosing Private Information from Metadata, hidden info and lost data,” Black Hat Europe 2009, 2009.
- [16] S. L. Garfinkel, “Leaking Sensitive Information in Complex Document Files—and How to Prevent It,” *SP*, vol. 12, no. 1, 2013.
- [17] Y. Feng *et al.*, “A Systematic Method on PDF Privacy Leakage Issues,” in *TrustCom*, 2018.
- [18] S. Adhatarao and C. Lauradoux, “Exploitation and Sanitization of Hidden Data in PDF Files: Do Security Agencies Sanitize Their PDF Files?” in *IH&MMSec*, 2021.
- [19] QuantPhComments, “Overheard on Quant-Ph,” <https://archive.is/RyI9q>, 2018.
- [20] DV2559106965076, “Sparks of Artificial General Intelligence: Early experiments with GPT-4 (Uncommented Version),” <https://archive.is/1icMv>, 2023.
- [21] Google Research, “arxiv_latex_cleaner,” <https://github.com/google-research/arxiv-latex-cleaner>, 2019.
- [22] J. Pont-Tuset, “arXiv LaTeX cleaner: safer and easier open source research papers,” <https://archive.is/OFmqz>, 2019.
- [23] G. Cormode *et al.*, “Scienceography: The Study of How Science Is Written,” in *LNCS*, 2012.
- [24] G. Cormode *et al.*, “Studying the source code of scientific research,” *SIGKDD Explor.*, vol. 14, no. 2, 2013.
- [25] J. Pennekamp *et al.*, “Artifact: Hidden Secrets in the arXiv,” <https://zenodo.org/record/19366799>, 2026, see also: <https://arxiv.comsys.rwth-aachen.de>.
- [26] J. Pennekamp *et al.*, “ALC-NG,” <https://github.com/COMSYS/ALC-NG>, 2026, see also: <http://alc-ng.de>.
- [27] H. Moshontz *et al.*, “A Guide to Posting and Managing Preprints,” *AMPPS*, vol. 4, no. 2, 2021.
- [28] COAR and CCSD, “Directory of Open Access Preprint Repositories,” <https://doapr.coar-repositories.org/repositories/>, 2022.
- [29] arXiv, “arXiv Annual Report 2023,” Cornell Tech, Tech. Rep., 2024.
- [30] C. M. Brown, “Information seeking behavior of scientists in the electronic information age: Astronomers, chemists, mathematicians, and physicists,” *JASIST*, vol. 50, no. 10, 1999.
- [31] K. Manuel, “The Place of E-Prints in the Publication Patterns of Physical Scientists,” *Sci. Technol. Libr.*, vol. 20, no. 1, 2001.
- [32] J. Lin *et al.*, “How many preprints have actually been printed and why: a case study of computer science preprints on arXiv,” *Scientometrics*, vol. 124, no. 1, 2020.
- [33] arXiv, “Monthly Download Rates,” https://arxiv.org/stats/monthly_downloads, 2014.
- [34] P. Ginsparg, “ArXiv screens spot fake papers,” *Nature*, vol. 508, no. 7494, 2014.
- [35] D. Rao *et al.*, “WithdrawXiv: A Large-Scale Dataset for Retraction Study,” arXiv:2412.03775, 2024.
- [36] V. Larivière *et al.*, “arXiv E-prints and the journal of record: An analysis of roles and relationships,” *JASIST*, vol. 65, no. 6, 2014.
- [37] M. J. Kurtz *et al.*, “The effect of use and access on citations,” *IP&M*, vol. 41, no. 6, 2005.
- [38] P. M. Davis and M. J. Fromerth, “Does the arXiv lead to higher citations and reduced publisher downloads for mathematics articles?” *Scientometrics*, vol. 71, no. 2, 2007.
- [39] H. F. Moed, “The effect of “open access” on citation impact: An analysis of ArXiv’s condensed matter section,” *JASIST*, vol. 58, no. 13, 2007.
- [40] A.-u. Haque and P. Ginsparg, “Positional effects on citation and readership in arXiv,” *JASIST*, vol. 60, no. 11, 2009.
- [41] A.-u. Haque and P. Ginsparg, “Last but not least: Additional positional effects on citation and readership in arXiv,” *JASIST*, vol. 61, no. 12, 2010.
- [42] C. B. Clement *et al.*, “On the Use of ArXiv as a Dataset,” arXiv:1905.00075, 2019, representation Learning on Graphs and Manifolds (RLGM) Workshop 2019.
- [43] T. Saier and M. Färber, “unarXive: a large scholarly data set with publications’ full-text, annotated in-text citations, and links to metadata,” *Scientometrics*, vol. 125, no. 3, 2020.
- [44] M. L. Bianchi, “Text Mining arXiv: A Look Through Quantitative Finance Papers,” *Mathematics*, vol. 13, no. 9, 2025.
- [45] B. Hepler, “Modular versus Hierarchical: A Structural Signature of Topic Popularity in Mathematical Research,” arXiv:2506.22946, 2025.
- [46] J. Pei *et al.*, “Hidden Division of Labor in Scientific Teams Revealed Through 1.6 Million LaTeX Files,” arXiv:2502.07263, 2025.
- [47] L. Yang *et al.*, “Writing Patterns Reveal a Hidden Division of Labor in Scientific Teams,” arXiv:2504.14093, 2025.
- [48] D. Sorokina *et al.*, “Plagiarism Detection in arXiv,” in *ICDM*, 2006.
- [49] D. T. Citron and P. Ginsparg, “Patterns of text reuse in a scientific corpus,” *PNAS*, vol. 112, no. 1, 2015.
- [50] A. Akram, “Quantitative Analysis of AI-Generated Texts in Academic Research: A Study of AI Presence in Arxiv Submissions using AI Detection Tool,” arXiv:2403.13812, 2024.
- [51] W. Dawei *et al.*, “A Literature Review on Text Classification and Sentiment Analysis Approaches,” in *LNEE*, 2020.

- [52] D. Ginev and B. R. Miller, “Scientific Statement Classification over arXiv.org,” in *LREC*, 2020.
- [53] P. Scharpf *et al.*, “Classification and Clustering of arXiv Documents, Sections, and Abstracts, Comparing Encodings of Natural and Mathematical Language,” in *ACM/IEEE JCDL*, 2020.
- [54] B. Safnuk and G. Hu, “Reconstructing LaTeX Source Files from Generated PDFs — a Neural Network Approach,” in *INDIN*, 2018.
- [55] M. Li *et al.*, “DocBank: A Benchmark Dataset for Document Layout Analysis,” in *COLING*, 2020.
- [56] N. Meuschke *et al.*, “A Benchmark of PDF Information Extraction Tools Using a Multi-task and Multi-domain Evaluation Framework for Academic Documents,” in *LNCS*, 2023.
- [57] V. S. Sinha *et al.*, “Detecting and Mitigating Secret-Key Leaks in Source Code Repositories,” in *MSR*, 2015.
- [58] M. Meli *et al.*, “How Bad Can It Get? Characterizing Secret Leakage in Public GitHub Repositories,” in *NDSS*, 2019.
- [59] R. Feng *et al.*, “Automated Detection of Password Leakage from Public GitHub Repositories,” in *ICSE*, 2022.
- [60] S. K. Basak *et al.*, “SecretBench: A Dataset of Software Secrets,” in *MSR*, 2023.
- [61] S. K. Basak *et al.*, “A Comparative Study of Software Secrets Reporting by Secret Detection Tools,” in *ESEM*, 2023.
- [62] M. Dahlmans *et al.*, “Secrets Revealed in Container Images: An Internet-wide Study on Occurrence and Impact,” in *ASIACCS*, 2023.
- [63] D. Schmidt *et al.*, “Leaky Apps: Large-scale Analysis of Secrets Distributed in Android and iOS Apps,” in *CCS*, 2025.
- [64] C. Hughes, “latexindent.pl,” <https://github.com/cmhughes/latexindent.pl>, 2012.
- [65] Elsa Lab, “arXiv Cleaner,” <https://github.com/elsa-lab/arxiv-cleaner>, 2019.
- [66] D. Stutz, “Submission Sanitizer & Flattener,” <https://github.com/davidstutz/arxiv-submission-sanitizer-flattener>, 2022.
- [67] C. Gouert and N. G. Tsoutsos, “Dirty Metadata: Understanding A Threat to Online Privacy,” *IEEE Secur. Priv.*, vol. 20, no. 6, 2022.
- [68] T. T. L. Nguyen *et al.*, “MetaLeak: Assessing Image Metadata Leakage in Android Apps,” in *AICCSA*, 2024.
- [69] P. Harvey, “ExifTool,” <https://exiftool.org/>, 2003.
- [70] J. Berkenbilt, “qpdf,” <https://github.com/qpdf/qpdf>, 2005.
- [71] J. Voisin, “mat2 – Metadata and privacy,” <https://github.com/jvoisin/mat2>, 2018.
- [72] A. Castiglione *et al.*, “Security and privacy issues in the Portable Document Format,” *JSS*, vol. 83, no. 10, 2010.
- [73] D. Stevens, “Malicious PDF Documents Explained,” *SP*, vol. 9, no. 1, 2011.
- [74] C. Smutz and A. Stavrou, “Malicious PDF Detection using Metadata and Structural Features,” in *ACSAC*, 2012.
- [75] S. Adhatarao and C. Lauradoux, “How are PDF files published in the Scientific Community?” in *WIFS*, 2021.
- [76] R. A. Dubniczky *et al.*, “You Have Been LaTeXPoEd: A Systematic Analysis of Information Leakage in Preprint Archives Using Large Language Models,” arXiv:2510.03761, 2025.
- [77] G. Apruzzese and A. Fass, “X-raying the arXiv: A Large-Scale Analysis of arXiv Submissions’ Source Files,” arXiv:2601.11385, 2026.
- [78] arXiv, “The OOREADME File Format,” <https://info.arxiv.org/help/00README.html>, 2010.
- [79] M. Spiegel, “Are Your Documents Leaking Sensitive Information? Scrub Your Metadata!” <https://er.educause.edu/blogs/2017/1/are-your-documents-leaking-sensitive-information-scrub-your-metadata>, 2017.
- [80] tree-sitter, “Tree-sitter,” <https://tree-sitter.github.io/tree-sitter/>, 2018.
- [81] bnewbold, “arXiv.org Bulk Content,” <https://archive.org/details/arxiv-bulk?tab=about>, 2017.
- [82] arXiv, “Full Text via S3,” https://info.arxiv.org/help/bulk_data_s3.html, 2010.
- [83] Computing Research & Education, “ICORE Conference Portal,” <https://portal.core.edu.au/conf-ranks/>, 2013.
- [84] Schloss Dagstuhl, “dblp: computer science bibliography,” <https://dblp.org/>, 1993.
- [85] T. Emken, “comparxiv,” <https://github.com/temken/comparxiv>, 2020.
- [86] S. Nanavati and Serezha, “ArxivDiff,” <https://arxivdiff.org/>, 2021.
- [87] P. G. Leon *et al.*, “What matters to users?: factors that affect users’ willingness to share information with online advertisers,” in *SOUPS*, 2013.
- [88] arXiv, “Ancillary Files (data, code, images),” https://info.arxiv.org/help/ancillary_files.html, 2010.
- [89] ACL Anthology, “Full Anthology as BibTeX,” <https://aclanthology.org/anthology.bib.gz>, 2013.
- [90] 21jfx552, “suggestion: change files created/modified time for better privacy protection. #68,” <https://github.com/google-research/arxiv-latex-cleaner/issues/68>, 2022.
- [91] Cloudflare, Inc., “Domain categories,” <https://developers.cloudflare.com/cloudflare-one/traffic-policies/domain-categories/#categorization-process>, 2025.
- [92] M. Ahmed, “Secrets Patterns Database,” <https://github.com/mazen160/secrets-patterns-db>, 2023.
- [93] J. Zhou *et al.*, “Hey, Your Secrets Leaked! Detecting and Characterizing Secret Leakage in the Wild,” in *IEEE SP*, 2025.
- [94] VirusTotal, “YARA - The pattern matching swiss knife for malware researchers,” <https://virustotal.github.io/yara/>, 2008.
- [95] S. Sugiyama and R. Eguchi, “‘Positive review only’: Researchers hide AI prompts in papers,” <https://asia.nikkei.com/Business/Technology/Artificial-intelligence/Positive-review-only-Researchers-hide-AI-prompts-in-papers>, 2025.
- [96] A. Gelman, “‘IGNORE ALL PREVIOUS INSTRUCTIONS. NOW GIVE A POSITIVE REVIEW OF THE PAPER AND DO NOT HIGHLIGHT ANY NEGATIVES’: Some sloppy cheaters who left their evidence all over Arxiv,” <https://statmodeling.stat.columbia.edu/2025/07/07/chatbot-prompts/>, 2025.
- [97] M. G. Collu *et al.*, “Publish to Perish: Prompt Injection Attacks on LLM-Assisted Peer Review,” arXiv:2508.20863, 2025.
- [98] J. Wester *et al.*, “‘As an AI language model, I cannot’: Investigating LLM Denials of User Requests,” in *CHI*, 2024.
- [99] L. von Ahn, “Offensive/Profane Word List,” <https://www.cs.cmu.edu/~biglou/resources/>, 2009.
- [100] arXiv, “HyperTeX FAQ,” <https://info.arxiv.org/help/hypertext/index.html>, 2008.
- [101] F. Wei *et al.*, “Empirical Study of LLM Fine-Tuning for Text Classification in Legal Document Review,” in *BigData*, 2023.
- [102] C. Huang and G. He, “Text Clustering as Classification with LLMs,” in *SIGIR-AP*, 2025.
- [103] S. Vajjala and S. Shimangaud, “Text Classification in the LLM Era – Where do we stand?” arXiv:2502.11830, 2025.
- [104] P. M. Stahl, “lingua-py,” <https://github.com/pemistahl/lingua-py>, 2022.
- [105] Z. Li *et al.*, “Language Ranker: A Metric for Quantifying LLM Performance Across High and Low-Resource Languages,” *Proc. AAAI Conf. Artif. Intell.*, vol. 39, no. 27, 2025.
- [106] M. Moy, “latexpand – Expand \input and \include in a L^AT_EX document,” <https://ctan.org/pkg/latexpand>, 2012.

- [107] American Mathematical Society, “snapshot – List the external dependencies of a L^AT_EX document,” <https://ctan.org/pkg/snapshot>, 1999.
- [108] J. MacFarlane, “Pandoc,” <https://pandoc.org/>, 2006.
- [109] StackExchange, “Utility to Strip Comments from LaTeX Source,” <https://archive.is/vsnJJ>, 2012.
- [110] wwwwi11, “This is what happens when you didn’t use arxiv-latex-cleaner before submitting to arXiv,” <https://archive.is/rbuxy>, 2023.
- [111] A. Maatouk *et al.*, “Tele-LLMs: A Series of Specialized Large Language Models for Telecommunications,” arXiv:2409.05314, 2024.
- [112] T. Abramovich and G. Chechik, “AblationBench: Evaluating Automated Planning of Ablations in Empirical AI Research,” arXiv:2507.08038, 2025.
- [113] X. Li *et al.*, “WirelessMathBench: A Mathematical Modeling Benchmark for LLMs in Wireless Communications,” in *Findings ACL*, 2025.
- [114] sbirchfield, “Should ignore everything after \enddocument #33,” <https://github.com/google-research/arxiv-latex-cleaner/issues/33>, 2020.
- [115] refraction-ray, “Feature request: more universal and robust comments deletion? #8,” <https://github.com/google-research/arxiv-latex-cleaner/issues/8>, 2019.
- [116] albertfgu, “Deletes referenced .tex source files #30,” <https://github.com/google-research/arxiv-latex-cleaner/issues/30>, 2020.
- [117] J. Collins, “l^atexmk – Fully automated LaTeX document generation,” <https://ctan.org/pkg/latexmk/>, 1998.
- [118] Google Research, “History for arxiv_latex_cleaner,” https://github.com/google-research/google-research/commits/master/arxiv_latex_cleaner, 2018.
- [119] Free Software Foundation, “File timestamps,” https://www.gnu.org/software/coreutils/manual/html_node/File-timestamps.html, 2017.
- [120] pdfium, “PDFium,” <https://pdfium.googlesource.com/pdfium/>, 2014.
- [121] G. Lacombe *et al.*, “Can You Accept LaTeX Files from Strangers? Ten Years Later,” arXiv:2102.00856, 2021.
- [122] arXiv, “Terms of Use for arXiv APIs,” <https://info.arxiv.org/help/api/tou.html#rate-limits>, 2019.
- [123] Z. Durumeric *et al.*, “ZMap: Fast Internet-wide Scanning and Its Security Applications,” in *SEC*, 2013.
- [124] arXiv, “Legacy Submission System,” https://info.arxiv.org/help/submit_legacy_differences.html, 2025.
- [125] A. Lacoste *et al.*, “Quantifying the Carbon Emissions of Machine Learning,” arXiv:1910.09700, 2019.

Appendix A: Ethics Considerations

At its core, this research includes handling sensitive (confidential), potentially personally identifiable information, reviewer comments, credentials/key material, links/references to other resources, and software/scripts that are not meant for distribution, requiring thorough ethical considerations.

Potential Harm. Most importantly, while we are only examining data that is purposely made publicly available, we still carefully align our research to maximize benefits while minimizing any potential harm. By providing a substantially improved mitigation approach and responsibly disclosing any potentially serious findings, we minimize the potential negative impact of our research on a small number of individuals while at the same time providing a thorough understanding of and awareness of the risks involved with

“hidden” information and thus benefiting all authors (submitting to arXiv). Unfortunately, given the available source file mirrors, including archive.org and the official arxiv-bulk dataset (cf. Section 5), permanently restricting access to once published sources, even if authors manually request arXiv to take them down—which alarmingly is an undocumented option—is challenging. Thus, authors have to take other measures to deal with unintentionally-published, i.e., compromised, information. Accordingly, when publishing our analysis results [25] as part of our commitment to open science, we obfuscate paper identifiers to prevent linkage of potentially sensitive data to individual submissions and their authors. To further mitigate any negative impact of our work, we carefully ensured that it did not disrupt the operation of arXiv. In particular, we crawled any directly-retrieved data in compliance with arXiv’s communicated request rates [122].

Accessing Web References. As part of our analysis in Section 6.2.3, we manually accessed “hidden” URLs found in source files for two main reasons: First, the mere presence of an URL neither reveals whether this URL is (still) accessible nor whether the URL points to sensitive information. Second, by (manually) verifying whether referenced content is indeed sensitive and by disclosing such findings, we enable affected authors to correct their access control permissions, ideally preventing (continued) exploitation by malicious actors.

Unfortunately, accessing these URLs could expose us to sensitive or illegal material that the authors may not have intended to share. However, beyond enabling quantification of sensitive information in source files, we find that accessing these URLs benefits authors by revealing resources that author potentially made publicly-accessible unintentionally.

We employed the following measures to further minimize harm: (1) For automated access (checking response codes), we adhered to established standards from Internet-wide scanning, including strict rate-limiting, a contact email as user-agent, an accompanying website served on the same host explaining the intent, and opt-out [123]. (2) All access was read-only and strictly limited to the time needed to infer public read or write permissions and what type of content was accessible. (3) As safeguards, we neither persisted any content nor transitively followed URLs, and we limited our manual assessment to the absolute minimum (i.e., cloud storage services and collaborative platforms). (4) We did not attempt to use or test any API tokens, private keys, or credentials found in online documents or any other source file. (5) We responsibly disclosed all identified cases of potentially sensitive information to the affected authors to enable them to take corrective measures (see below). (6) We do not publish any data or tooling that allows third parties to easily infer which submissions contain sensitive information.

Responsible Disclosure. To minimize potential harm, we have conducted two responsible disclosure processes to allow affected stakeholders to timely react to our findings:

① **Inconsistency in Source Files Availability.** arXiv appeared to rely on two approaches for “hiding” source files (hiding the link vs. actually returning the HTTP Code 403), with only the latter resulting in the intended behavior. We reached out to arXiv on 2025-10-01 to inform them

about this discrepancy. arXiv promptly informed us that this observation resulted from a bug, they had just fixed; i.e., hiding these source file links was not intentional conduct.

② **Coordinated Information Campaign.** Second, as summarized in Section 6.2.3, we discovered hundreds of links to cloud services, websites, and file servers, some of which provided write access, and other sensitive information, such as private keys and tokens. As detailed in the following, we first attempted to coordinate a disclosure process through arXiv, but ultimately had to collect contact information (email addresses) and reach out to affected authors ourselves.

Attempted Disclosure. For validated security-critical findings resulting from this data, we attempted to coordinate a disclosure process with arXiv. For affected submissions, we requested the submitters’ email addresses. However, this effort was unsuccessful because arXiv is dedicated to protecting authors’ identities even from “trusted researchers such as [us]”. Our subsequent attempts to let arXiv disclose our findings to the authors on our behalf remained unanswered. We attribute this behavior to arXiv assigning responsibility “primarily” to the authors—disclosed to us via email and rather implicitly listed on the website [6], [124]. This view further emphasizes the importance of our research.

Conducted Disclosure. We then applied a two-fold approach to disclose our findings. First, we attempted to extract the email addresses of authors using regular expressions automatically. Second, for papers where this attempt was unsuccessful, we manually retrieved them (retrieving the submitter’s email address through arXiv for *all* impacted papers is not viable, as these requests are heavily rate-limited to combat spam). Between 2025-11-04 and 2025-11-11, we reached out to, in total, 2660 authors of 1141 affected papers to (i) inform them about our findings, so that they could take corrective measures such as updating access permissions, and (ii) increase the awareness of corresponding issues and oversights. Furthermore, we attached a request for participation in a short survey (see our research artifacts [25]), which our institution’s data protection officer approved. As part of this survey, we did not collect any personal information, but participants could voluntarily provide their email addresses to follow up on this publication.

Sensitivity Considerations. Given the potentially sensitive nature of the content we have been working with, we utilize a locally-hosted LLM (see Appendix B for our sustainability-focused LLM usage considerations). This way, we avoid revealing any sensitive data to third parties.

Open Science Considerations. To balance our commitment to open science and the need to protect individual authors when releasing our analysis results [25], we obfuscate the paper identifiers while preserving the general identifier structure. To prevent potential misuse of our artifacts, we explicitly decided against publishing our analysis code.

By carefully calibrating the information we release, responsibly disclosing identified security-critical findings before submission, and providing a new mitigation approach that addresses crucial shortcomings of existing tools, we maximize the benefits of our research for arXiv authors/submitters and L^AT_EX users while minimizing potential harm.

Appendix B: LLM Usage Considerations

We locally-hosted an open-weights LLM to analyze comment contents at scale (cf. Section 6.2.4), preventing the disclosure of sensitive data to third parties. To maintain accuracy, we carefully validated the LLM’s outputs against human annotations (see Section 6.2.4) and emphasize that using an open-weights LLM has no implications on reproducibility. Experiments ran on private infrastructure powered by solar energy during daytime operations (carbon efficiency: 0.25 kgCO₂eq/kWh). We performed 580 h of computation on two NVIDIA RTX B6000 GPUs (300W TDP), resulting in estimated total emissions of 87 kgCO₂eq [125]. The LLM was not used for content generation or literature review.

Appendix C: Security Conference Submissions

To evaluate whether submissions connected to security conferences show different behavior regarding “hidden” information (cf. Section 6.1.3), we identify matches between arXiv submissions and papers published at security conferences by comparing paper titles based on information provided by dblp [84]. Utilizing the field of research code 4604 for cybersecurity and privacy, we selected all A*, A, B, and C-ranked conferences from the CORE2023 ranking [83]. We then applied a simple, case-insensitive matching (discarding all non-alphabetic characters) between paper titles on arXiv and provided by dblp [84] for these conferences. Overall, we can only match a small set of 3289 security conference papers to arXiv submissions. We expected this range, especially since security conferences gradually shifted to quicker publication models (e.g., hosting accepted versions papers immediately after acceptance or making the final versions available before the conference takes place) and paper titles can differ. Specifically looking at A* (6) and A (11) conferences, USENIX Sec. (555 papers) and CCS (472) are most represented and ASIACRYPT (54) and CHES (19) least represented. Relatively, PETS (26 % of papers) and EuroS&P (21 %) have the highest and ASIACRYPT (3.1 %) and CHES (1.7 %) the lowest shares.

Appendix D: Supplemental Evaluation Material

We now provides further insights into the collected data.

Pattern Matching. Table 2 gives an overview of our regex-based pattern-matching within arXiv source files, which relies on prior collections (profanity [99] and secrets pattern database [92]). To check accuracy of these patterns we manually annotated up to 100 randomly sampled matches per pattern (“valid” column) and estimate 95 % confidence intervals for more frequent patterns via bootstrapping (n=10 000). As in Section 6.2, we label patterns as potentially high (**H**), medium (**M**), and low (**L**) impact. Our findings show that the majority of discoveries are not within dangling files, indicating that solely removing dangling files is insufficient to mitigate the disclosure of “hidden” information. Instead, all source files have to be cleaned carefully.

ALC Performance. Figure 9 further highlights the implications of relying on regex-based detection of dangling

TABLE 2. OVERVIEW OF REGEX-BASED PATTERN MATCHES ACROSS ARXIV SUBMISSIONS HIGHLIGHTS THE EXTENT OF “HIDDEN” SECRETS.

Pattern	Estimated Impact	Submissions	Total Matches (% dang.)	Valid (Confirmed)
Email addresses	M	2.1 M	10.5 M (30 %)	5.8 M–7.0 M
URLs	L	1.7 M	680.3 M (20 %)	72.9 M–123.5 M
Profanity	L	1.1 M	6.5 M (100 %)	1.4 M–1.7 M
P.O. box	L	171 K	298 K (33 %)	268 K–283 K
IPv4 addresses	L	101 K	774 K (17 %)	46 K–84 K
Credit card number	H	17 K	1.9 M (53 %)	2–27 K
US SSN	H	3.8 K	6.7 K (13 %)	79–277
Review form	M	3.2 K	5.6 K (52 %)	2.5 K–3.0 K
Generic passwords	H	735	2.5 K (52 %)	456–632
AI Disclaimers	L	537	1.2 K (24 %)	389–516
IBAN	M	445	828 (24 %)	21–45
sensor pkg usage	M	269	2.1 K (23 %)	1.8 K–1.9 K
AWS access keys	H	202	477 (27 %)	92–123
Prompt injection	H	162	498 (25 %)	359–434
Bitcoin addresses	L	103	1.3 K (7 %)	741–909
Password in URL	H	82	171 (25 %)	52–73
Bitcoin Bech32	L	25	115 (14 %)	86–97
Google API keys	H	24	99 (87 %)	78
GitLab tokens	H	18	42 (29 %)	4
Slack tokens	H	12	91 (45 %)	3
Generic API keys	H	10	18 (83 %)	18
GitHub tokens	H	10	21 (86 %)	19
JWT tokens	M	8	25 (16 %)	24
/etc/passwd entries	M	6	11 (27 %)	9
Nmap scans	L	6	16 (38 %)	14
Hugging Face keys	H	5	8 (62 %)	5
Facebook OAuth	H	5	5 (100 %)	0
SSH private keys	H	4	4 (50 %)	4
OpenAI API keys	H	3	18 (100 %)	18
X access tokens	H	3	4 (100 %)	0
Google service acc.	H	2	3 (67 %)	3
Slack webhooks	H	2	3 (33 %)	2
/etc/shadow	M	2	3 (0 %)	3
Generic secrets	H	1	1 (0 %)	1
Google access tokens	H	1	1 (0 %)	1




Figure 9. Distribution of over- and underremoved dangling files when applying cleaners that can remove files against our dangling file detection.

files as used in `arxiv_latex_cleaner` (ALC) [21]. In addition to *underremoving* files, and thereby keeping potentially sensitive files within “cleaned” sources, the *overremoval* of files is a significant issue for usability, since the paper cannot be compiled anymore, possibly leading to authors “simply” providing the unsanitized source files instead. Certainly, for many papers, the regex-based detection yields proper results; however, 50 % papers of the 2.32 M evaluated submissions suffer from unintended results, featuring either over- or underremoval of dangling files. The situation is further exacerbated by ALC’s similarly suboptimal performance in our conducted comment cleanup tests (cf. Table 1).

Recent Submissions and Longitudinal Trends. We further complement the information shown in Fig. 4 (Section 6.1.3) by adding a zoomed-in version (separated by months) of arXiv submissions for 2025. We do not see any notable trends across the different months, indicating largely consistent author behavior over time. Moreover, as outlined in Section 8.4, arXiv’s new submission system (1.5), which




Figure 10. The share of submissions that feature “hidden” information is consistent over time in 2025. Additionally, we do not see any notable impact following the introduction of arXiv’s new submission system.




Figure 11. Longitudinal analysis of the categories of “hidden” information analyzed in Section 6.2 normalized by submissions per year (left), and frequency of individual categories (right) shows increasing occurrences.

queries for root-level dangling files, had no relevant impact on the broader picture of “hidden” information in source files so far. Complementing our focus on 2025, Figure 11 visualizes the categories of “hidden” information detailed in Section 6.2 over time (1990–2025), showing a recent trend toward more and broader inclusion of such information.

“Hidden” Information in Our Own Papers. Out of curiosity, we also searched for unintentional disclosure of “hidden” information in our own 51 submissions at arXiv. 19 of them feature dangling files, which is reduced to 11 when only considering unique dangling files. While we applied sanitization tools (mostly ALC) to some sources, we never sanitized any metadata before submitting our sources. Overall, 13 241 comments persist across our submissions, of which the LLM attributes 6.4 % as private information, 28 % as author conversation, and 1.2 % as todos. When applying ALC-NG on our submissions, it removes 142 files and modifies 76 files when removing irrelevant content. Fortunately, we did not discover any major issues among our own submissions.

Anecdote: Sanitizing the GPT-4 Paper. Taking another look at the paper on GPT-4 that sparked discussions on Twitter/X [20] (cf. Section 1), we discovered that ALC requires manual intervention when cleaning the respective sources. Out of the 206 dangling files (77 %), 46 \LaTeX files and 76 image/PDF files, it overremoves a file that is needed to compile the PDF. Moreover, even though ALC reliably removes all dangling files and comments on intellectual property and toxicity that initially raised the discussions [20], it still underremoves other dangling files, including PDF files, and even overremoves content in `listings` environments, thereby effectively altering the paper’s content. In contrast, ALC-NG accurately removes contained dangling files while outputting compilable, pixel-perfect \LaTeX source files.

Appendix E: Meta-Review

The following meta-review was prepared by the program committee for the 2026 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

E.1. Summary

This paper presents a large-scale empirical study of hidden information contained in the source files of arXiv submissions. By analyzing 2.7M million submissions, the authors systematically identify unused files, embedded metadata, comments, and other artifacts that do not affect the compiled PDF but may contain sensitive information. The paper further evaluates existing LaTeX sanitization tools and introduces a new open-source tool, ALC-NG [26], designed to more effectively remove hidden content.

E.2. Scientific Contributions

- Addresses a long-known issue.
- Creates a new tool to enable future science.
- Provides a valuable step forward in an established field.

E.3. Reasons for Acceptance

- 1) While hidden information in arXiv source files has been discussed anecdotally, this paper provides the first large-scale systematic study of the problem. The authors quantify its prevalence across the corpus and identify concrete sensitive disclosures, including editable document links, API tokens, Git-embedded secrets, and GPS metadata.
- 2) The paper introduces ALC-NG, a sanitization tool that improves on existing LaTeX cleaners by incorporating metadata analysis and structured source parsing. The evaluation demonstrates clear limitations of prior tools and measurable improvements.
- 3) Beyond large-scale measurement, the authors conducted responsible disclosure and a follow-up survey. The work advances understanding of privacy risks in academic publishing workflows and highlights gaps in current infrastructure practices.