SoK: Anti-Forensics Evolution, Platform Coverage, and Mobile Residual Evidence

Uju Okoye[a,b], Roohana Karim[a,b], Abdur Onik[a,b], Ibrahim Baggili[a,b]

[a] *Baggili(i) Truth (BiT) Lab, Center of Computation & Technology, Baton Rouge, LA, USA*
[b] *Division of Computer Science & Engineering, Louisiana State University, Baton Rouge, LA, USA*

Abstract

This Systematization of Knowledge (SoK) analyzes 197 primary studies (2005–2024) to examine how anti-forensics (AF) techniques have evolved, their distribution across platforms, and what forensic artifacts persist after mobile AF activity. We report three key findings. First, AF has matured from traditional data hiding and artifact wiping toward deception-oriented approaches; data fabrication now dominates, and AI-driven anti-forensics (AAF) has emerged but exclusively in multimedia pipelines, with no mobile validation. Second, platform coverage is imbalanced: Windows and Android dominate, while iOS and macOS each appear in only a few studies, and cross-platform validation is rare. Third, mobile AF tends to displace rather than erase evidence: across 18 studies, 93 residual observations persist in SQLite/WAL, logs, caches, and filesystem metadata, concentrated in app-private (51.6%) and OS-level (29.0%) storage. This SoK contributes a new two-part taxonomy linking AF strategy families (with an AAF extension) to mobile residual evidence by store type and location, along with an empirical baseline for cross-platform AF research. The findings underscore the need for standardized benchmarks, expanded iOS coverage, and mobile-focused AAF evaluation to strengthen forensic reliability.

Keywords: mobile anti-forensics, digital forensics, anti-forensics techniques, residual artifacts, Android, iOS

1. Introduction

In September 2024, an international law enforcement coalition dismantled "Ghost," an encrypted messaging platform used by organized crime networks across multiple continents.[1] Built on modified Android devices, Ghost featured self-destructing messages, remote wipe capabilities, and hardware modifications designed to defeat forensic acquisition. Despite these measures, investigators recovered critical evidence from seized devices showing that even sophisticated anti-forensic architectures do not guarantee evidentiary destruction.

At the same time, investigations involving ephemeral messaging and deletion features highlight that evidentiary loss can be difficult to validate without systematic knowledge of what should persist and where to look. Research on ephemeral messaging applications suggests such platforms frequently leave forensic traces in database logs, cached metadata, and system notifications (Heath et al., 2023), raising the question of whether investigators lacked the systematic knowledge to locate what remained. These different realities reveal a critical gap in digital forensic knowledge.

Anti-forensics encompasses deliberate actions intended to hinder, manipulate, or prevent the discovery of digital evidence (Harris, 2006; Rogers, 2006). Since early taxonomies identifying data hiding, artifact wiping, trail obfuscation, and attacks against forensic processes (Rogers, 2006), the field has expanded across desktops, servers, cloud, IoT, and mobile platforms, each with different forensic challenges. Existing surveys provide broad classifications (Conlan et al., 2016; Gül and Kugu, 2017; Neale, 2023), yet knowledge remains scattered. In particular, AI-driven anti-forensics (AAF), which uses generative models and adversarial examples to evade detection, has outpaced classical taxonomies, and previous surveys rarely quantify platform distribution, temporal trends, or validation beyond single-environment studies.

Within this landscape, mobile devices represent a critical yet understudied domain. Smartphones have become primary repositories of personal, transactional, and locational data, making them high-value targets in criminal investigations (Maček et al., 2016; DOrazio et al., 2014). Mobile anti-forensics (MAF) exploits platform-specific features like sandboxing, and proprietary filesystems (Distefano et al., 2010). Yet Anti-forensics (AF) research largely focuses on desktop environments, which leaves mobile-specific behaviors inadequately systematized. Importantly, studies on what forensic artifacts *persist* after MAF activity remain scattered across isolated case studies rather than synthesized into actionable guidance.

To address these gaps, we conduct a Systematization of Knowledge (SoK) of AF research published between 2005 and 2024, guided by three research questions:

Email addresses: uokoye4@lsu.edu (Uju Okoye), rkarim3@lsu.edu (Roohana Karim), aonik1@lsu.edu (Abdur Onik), baggili@gmail.com (Ibrahim Baggili)

[1] https://en.wikipedia.org/wiki/Ghost_(communication_network)

- **RQ1**: What anti-forensic techniques have been examined in the literature, and how have they evolved?
- **RQ2**: How is AF research distributed across platforms, and to what extent do studies address mobile devices?
- **RQ3**: What forensic artifacts persist after AF activity on mobile devices, and how can they be categorized?

We systematically identified and analyzed 197 primary studies from IEEE, Elsevier, ACM, and Springer venues and other scholarly venues. Each study was coded for: (i) the AF technique family, mapped to the taxonomy of Conlan et al. (2016); (ii) targeted platform(s) and device type(s); and (iii) for mobile-focused studies, residual forensic artifacts reported. This work makes three contributions:

- An empirical analysis of 197 studies spanning 2005 to 2024, showing how AF techniques, platforms, and domains have evolved over two decades.
- An extended taxonomy introducing *AAF* to capture adversarial machine learning, generative forgeries, and trace-suppression techniques absent from previous frameworks.
- A systematic characterization of residual artifacts persisting after mobile AF activity across SQLite databases, write-ahead logs, system logs, and filesystem metadata.

The remainder of this paper is organized as follows. Section 2 reviews prior AF surveys and mobile forensic studies. Section 3 details our systematic review methodology. Section 4 presents our taxonomy. Section 5 presents our results organized by research question, which are then discussed in Section 6. The limitations of this work are outlined in Section 7. Section 8 concludes with recommendations for future research.

2. Related Work

This section reviews prior research on anti-forensics (AF), tracing how the field evolved through definitions, taxonomies, surveys, and mobile experiments, and identifying gaps specific to mobile devices and residual evidence.

2.1. Overview of Anti-Forensics

Early definitions framed AF as attempts to compromise the availability, integrity, or usefulness of evidence for forensic investigation (Harris, 2006; Kessler, 2007). Rogers (2006) proposed a foundational categorization identifying four groups: *data hiding*, *artifact wiping*, *trail obfuscation*, and *attacks against forensic tools*. This taxonomy became canonical and remains widely adopted.

Subsequent work expanded both scope and granularity. Conlan et al. (2016) catalogued 308 tools and introduced a fifth category, *Possible Indications of Anti-Digital Forensics*, recognizing that everyday utilities (encryption apps, remote wiping services) can serve anti-forensic purposes even when not designed for them. Hasanabadi et al. (2020) applied game-theoretic modeling to characterize adversarial dynamics between investigators and anti-forensic actors, while Neale (2023) synthesized two decades of AF techniques and countermeasures, noting the continued immaturity of detection solutions.

Despite this progress, existing surveys share a common limitation: they were developed primarily from desktop and server tool analyses, with little systematic attention to how AF techniques distribute across computing platforms or how categories have evolved over time. Mobile devices, when mentioned, are treated briefly.

Research Gap (RG1). Existing AF taxonomies derive from desktop tool analyses and do not quantify cross-platform distribution or temporal evolution, leaving practitioners without empirical baselines.

2.2. Mobile-Focused AF Surveys and Studies

While general AF surveys established broad taxonomies, several studies have examined mobile-specific contexts, though coverage remains fragmented and platform-centric.

Early Android research demonstrated feasibility of mobile AF. Distefano et al. (2010) showed local paradigms for hiding data within Android filesystems; Albano et al. (2011b) demonstrated selective evidence erasure without cryptographic methods; and Karlsson and Glisson (2014) modified CyanogenMod to resist forensic extraction. On iOS, DOrazio et al. (2014) introduced techniques exploiting platform-specific backup mechanisms to conceal, delete, and insert data without permanent system alterations.

Subsequent surveys attempted consolidation. Maček et al. (2016) categorized Android forensic and anti-forensic techniques, while Barmpatsalou et al. (2013, 2018) surveyed mobile forensics broadly, treating AF primarily as an acquisition challenge. More recent work focuses on detection: Bhushan and Florance (2022) provided an overview of anti-forensic issues in Android forensics and proposed an SVM-based approach for identifying tampered encoded files, and Park et al. (2018) addressed vendor-specific vulnerabilities in Samsung encrypted backups.

These contributions share a common limitation: they remain siloed by platform, cataloguing tools or demonstrating feasibility rather than providing unified cross-platform analysis.

Research Gap (RG2). Mobile AF studies are predominantly Android-focused and tool-centric, with limited iOS coverage and no systematic cross-platform comparison.

Table 1: Summary of prior anti-forensics surveys, reviews, and selected artifact persistence studies (chronological). A checkmark (✓) indicates coverage; a cross (✗) denotes its absence.

Reference	Platform Focus	Type of Paper	General	Mobile	Artifacts	Summary
Distefano et al. (2010)	Android	Experimental Review	✗	✓	✗	Early Android AF study using local paradigms to hide/substitute data.
DOrazio et al. (2014)	iOS	Case Study	✗	✓	✗	iOS AF methods for concealing, deleting, and inserting data.
Anglano (2014)	Android	Experimental	✗	✗	✓	Recovery of WhatsApp SQLite DB and logs after deletion.
Conlan et al. (2016)	Cross-platform	SoK/Taxonomy	✓	✗	✗	Extended taxonomy of 300+ AF tools; minimal mobile focus.
Maček et al. (2016)	Android	Survey	✗	✓	✗	Survey of Android forensic and AF techniques.
Gül and Kugu (2017)	General	Survey	✓	✗	✗	Broad AF survey covering classical and modern AF methods.
Barmpatsalou et al. (2018)	Mixed mobile	Survey	✗	✓	✗	MDF survey; mentions AF as a challenge but not categorized.
Hasanabadi et al. (2020)	General	Survey	✓	✗	✗	Game-theoretic survey of AF, framing adversarial dynamics.
Salamh et al. (2020)	Android	Experimental	✗	✗	✓	WAL files retain deleted WhatsApp messages missed by tools.
Bhushan and Florance (2022)	Android	Overview	✗	✓	✗	Tool-based overview of Android AF issues.
Fukami et al. (2022)	Android	Experimental	✗	✗	✓	eMMC flash evaluation shows recoverable remnants after wiping.
Neale (2023)	General	Systematic Review	✓	✗	✗	Systematic review of AF countermeasures; immaturity of solutions noted.
Heath et al. (2023)	Cross-platform	Experimental	✗	✗	✓	Ephemeral messaging apps retain traces despite disappearing-message claims.
Kim et al. (2024)	Android	Experimental	✗	✗	✓	Wiping apps leave residual data in caches and Filesystem (FS) metadata.
Our SoK (2024)	Cross-platform	Systematization of Knowledge Review	✓	✓	✓	First unified review covering AF taxonomies, mobile AF, and residual artifacts.

2.3. Artifact Persistence and Residual Evidence

A recurring theme across AF studies is that data traces persist despite deletion or obfuscation attempts. While not surveys per se, these works collectively demonstrate the limits of anti-forensic methods and identify recovery opportunities for investigators.

Evidence of persistence has been repeatedly observed in database containers. Anglano (2014) reconstructed complete WhatsApp conversations from SQLite databases and associated log files, even after message deletion. Similarly, Salamh et al. (2020) showed that deleted content remained in SQLite Write-Ahead Log (WAL) files, escaping detection by common forensic tools.

Volatile and flash memory also retain sensitive data. Thing et al. (2010) developed a live memory forensic approach for Android that extracted browser and application traces missed by filesystem analysis. At the hardware level, Fukami et al. (2022) demonstrated that eMMC flash storage contains recoverable remnants after "secure" deletion.

Studies of wiping utilities confirm incomplete erasure. Kim et al. (2024) evaluated Android wiping applications and found residual data in caches and filesystem metadata. Abozaid et al. (2023) analyzed an Android shredding tool, showing high removal rates but also exposing limits of forensic validation tools.

Persistence extends to backups and cloud environments. Heath et al. (2023) analyzed ephemeral messaging applications (WhatsApp, Snapchat, Telegram), concluding that disappearing messages remain retrievable depending on implementation and extraction method. Park et al. (2018) demonstrated that Samsung encrypted backup data, which stores nearly all device information, can be decrypted through PIN recovery, showing that backup containers preserve artifacts that may no longer be accessible on the device itself.

These studies collectively show that AF rarely achieves complete erasure. Residual artifacts vary by technique, storage container, and acquisition method. However, findings remain scattered across isolated case studies. Table 1 summarizes prior work and scope. No review synthesizes these findings into a systematic characterization of what persists, where, and under what conditions.

Research Gap (RG3). Studies show artifacts persist in SQLite databases, WAL files, and caches, but no review systematically maps *which* artifacts survive specific AF techniques or *where* they remain across mobile platforms.

3. Methodology

We conducted a systematization of knowledge (SoK) to identify, evaluate, and synthesize existing research on AF techniques, their implementation on mobile devices, and the residual forensic artifacts that persist on mobile platforms despite AF activity. Our methodology follows the guidelines established by Kitchenham et al. (2004), which

provide a structured framework for transparent and reproducible evidence synthesis. We adopted a four-stage process consisting of: (i) planning and search strategy, (ii) data collection, (iii) screening and filtering, and (iv) data extraction and analysis.

3.1. Search Strategy

We initiated our literature search using keyword queries across multiple academic databases, including *IEEE Xplore*, *ScienceDirect*, *Scopus*, and *Google Scholar*. We conducted pilot searches to refine our terminology and assess coverage across these databases. These pilots indicated that Scopus provided the broadest and most consistent coverage of studies relevant to our research questions on anti-forensics, and we therefore selected it as our primary database. We designed three complementary queries to capture studies addressing: (i) anti-forensic techniques and tools, (ii) residual or forensic artifacts, and (iii) taxonomies or classifications of AF methods.

We applied these queries to titles, abstracts, and author keywords using Boolean operators to accommodate terminological variations (e.g., *"anti-forensics," "anti-forensic," "mobile forensics," "forensic artifacts"*). Table 2 presents the complete search strings and their result counts. We executed the search in July 2025 and restricted our scope to English-language, peer-reviewed studies published between 2005 and December 2024.

We selected 2005 as our start date because it marks the emergence of anti-forensics as a distinct research topic in the literature. We excluded 2025 publications to avoid partial-year bias and indexing lag. This search yielded an initial corpus of 798 records.

3.2. Data Collection

We exported all search results for systematic processing and cleaning. We removed duplicate records, book chapters, and short texts that did not constitute full research contributions (e.g., abstracts, posters, editorials). We also excluded studies for which full-text access was unavailable through institutional subscriptions. This initial cleaning reduced the corpus to 555 records.

To supplement our database search, we applied forward and backward snowballing using *Google Scholar* to identify additional relevant works that may have been missed by our initial queries. We screened all newly identified records using our inclusion and exclusion criteria. This process identified 11 candidate studies, of which 7 met all criteria and were added to the corpus.

3.3. Screening and Filtering

We conducted screening in two stages. In the first stage, we screened 555 articles by examining titles, abstracts, and keywords to identify studies addressing anti-forensic techniques, mobile platforms, or residual forensic artifacts. This stage excluded 315 records that fell outside our research scope, that is, studies that are not explicitly on AF, leaving 240 for full-text review.

In the second stage, we manually reviewed all full texts, examining introductions, research questions, contributions, methodologies, and conclusions. This stage excluded 50 additional records. Combined with 7 studies from snowballing, our final corpus comprised 197 primary studies.

Table 2: Search queries and results for the systematic review

Source	Search Query	Results
Scopus	TITLE-ABS-KEY(("anti-forensics" OR "anti forensic" OR "anti-forensic tools" OR "anti-forensic techniques" OR "forensic countermeasures" OR "forensic evasion") AND ("digital forensics" OR "mobile forensics"))	373
Scopus	TITLE-ABS-KEY(("forensic artifact*" OR "residual data" OR "data remnant*" OR "forensic trace*" OR "log trace*" OR "deleted data" OR "residual evidence" OR "forensic footprint*" OR "artifact recovery" OR "data carving") AND ("digital forensics" OR "mobile forensics"))	344
Scopus	TITLE-ABS-KEY(("anti-forensics" OR "anti forensic" OR "anti-forensic techniques" OR "forensic countermeasures") AND ("taxonomy" OR "classification" OR "categorization" OR "framework") AND ("digital forensics" OR "mobile forensics"))	81
Total		**798**

3.4. Data Extraction and Analysis

From the final corpus of 197 studies, we manually extracted data using a structured coding form. For each study, we recorded: (i) the AF techniques or tools reported, (ii) the targeted platform(s) and device type(s), and (iii) any residual forensic artifacts persisting on mobile devices following AF activity.

Because terminology varied substantially across the literature, we applied systematic normalization. We identified approximately 938 raw mentions of AF techniques, encompassing product names (e.g., BleachBit), operational actions (e.g., "secure-delete"), and descriptive phrases (e.g., "disappearing messages"). We mapped each mention to a subcategory within the taxonomy of Conlan et al. (2016): product names were mapped to their core operation (*e.g.*, `BleachBit` → Metadata wiping); synonymous terms were consolidated under a single label (*e.g.*, `sdelete`, `shred`, `secure-delete` → File wiping); and modern generative and adversarial manipulations such as GAN forgeries and diffusion-based inpainting were grouped under an extended AAF category that we introduced to capture emerging behaviors not covered by the original taxonomy. The complete mapping is provided in our supplementary repository.[2]

4. Taxonomy

This SoK separates *what adversaries do* (anti-forensic strategies and operations), *where they do it* (platforms), and *what remains* (residual artifacts). This separation makes diverse studies comparable across years, devices, and acquisition settings.

[2] https://github.com/mobile-af-artifacts/Mobile-Anti-Forensics-2025.git

4.1. Strategy-first anti-forensics taxonomy

Building on the strategy-first view introduced by Rogers (2006) and the extended, tool-grounded taxonomy of Conlan et al. (2016), we organize AF into *strategy families*: *Data Hiding (DH)*, *Artifact Wiping (AW)*, *Trail Obfuscation (TO)*, *Attacks Against Forensic Tools and Methods (AT)*, and *Possible Indications of Anti-Digital Forensics (PIA)* (Conlan et al., 2016). We retain *PIA* as a first-class category because it captures a distinct investigative question (*signals that AF may have occurred*) rather than an attacker operation, and it is frequently discussed in reviews but rarely validated with extractable indicators in mobile settings (Neale, 2023).

AAF extension. Recent work uses AI and ML as the *mechanism* of anti-forensics, often targeting *authenticity verification* (and forensic inference) rather than only *recoverability*. To preserve this signal during synthesis, we introduce a conservative sixth category: *AI-driven Anti-Forensics (AAF)*. We apply *AAF* only when the *adversary* uses AI and ML as an essential component of the anti-forensic method (e.g., generative alteration or inpainting, adversarial optimization against forensic detectors/provenance models, or learned suppression of manipulation traces), not when ML is used solely by the defender for detection or classification (Wu et al., 2020; Li et al., 2021; Liu et al., 2023). For traceability, we further code AAF into three operational subtypes: *Generation* (synthetic content creation or replacement) (Chen et al., 2019; Peng et al., 2022), *Evasion* (detector-targeted adversarial attacks) (Ding et al., 2021; Wang et al., 2021; Li et al., 2021; Fang and Stamm, 2023), and *Trace suppression* (learned removal/attenuation of forensic traces while preserving perceptual quality)(Liu et al., 2023; Wu and Sun, 2021).

Categories vs. sub-techniques. Strategy families are stable units for cross-year comparison, but they are too coarse to capture operational change. We therefore distinguish: (i) *Categories* (DH/AW/TO/AT/PIA/AAF), and (ii) *Sub-techniques* (concrete operations, such as file wiping, log manipulation, or adversarial example evasion). Figure 1 visualizes the strategy-family taxonomy used for normalization in this SoK and lists representative sub-techniques observed in the mobile subset of the reviewed corpus (*AAF* was not observed in the mobile subset).

4.2. A Taxonomy for Mobile Residual Evidence

To answer RQ3, we require a second taxonomy: not of attacker behavior, but of *residual evidence* after mobile AF activity. Prior work shows that mobile traces frequently persist in structured stores (e.g., SQLite and WAL), system/app logs, caches, and filesystem metadata, despite deletion or concealment attempts (Anglano, 2014; Salamh et al., 2020; Heath et al., 2023; Kim et al., 2024; Fukami et al., 2022). We therefore code each residual artifact observation along two orthogonal dimensions:

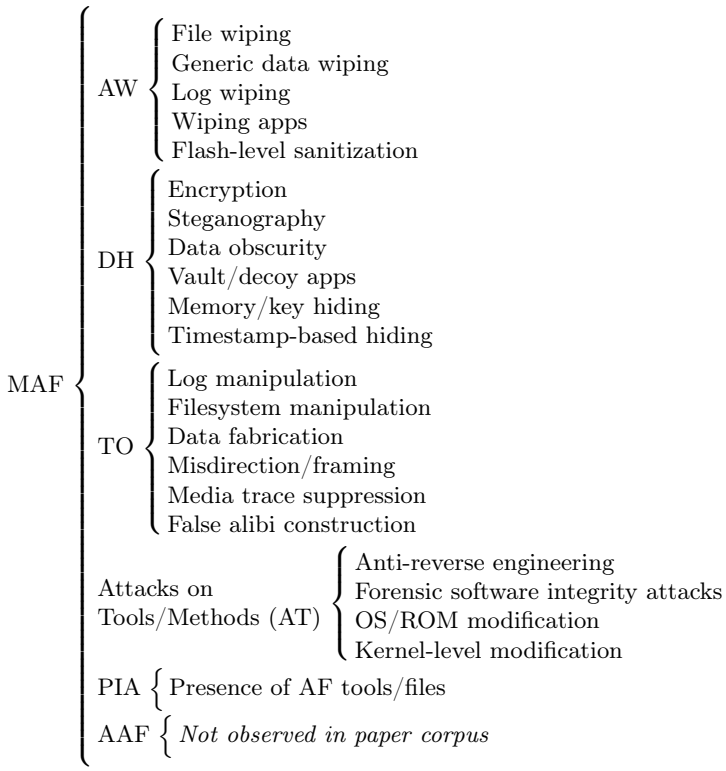

Figure 1: Strategy-first taxonomy used in this SoK for MAF. Top-level categories follow Conlan et al. (2016) (grounded in Rogers (2006)) with a conservative AAF extension for AI/ML-enabled AF (Wu et al., 2020; Li et al., 2021; Liu et al., 2023). Leaf nodes list representative sub-techniques observed in the mobile subset of the reviewed corpus (AAF not observed in the mobile subset).

Residual store type. We use six recurring store types: *SQLite (including WAL when applicable)*, *Logs*, *Caches and Thumbnails*, *Filesystem metadata*, *XML/plist*, and *Other* (for low-level remnants or artifacts that do not fit the above categories).

Residual location bucket. We categorize location as *App-private (sandbox)*, *OS-level*, *External or Shared*, and *Backup and Cloud* when reported. This aligns with practical acquisition constraints and prior mobile evidence studies that surface backup-related artifacts in vendor-specific ecosystems (DOrazio et al., 2014; Park et al., 2018).

Together, the strategy taxonomy (what is attempted) and the residual taxonomy (what persists) support a traceable synthesis: RQ1 quantifies categories and sub-techniques, RQ2 attaches platform coverage to the same categories, and RQ3 maps mobile AF categories to residual store types and locations.

5. Results

This section answers each research question using the taxonomy as the organizing lens. We emphasize three counting units: (i) *paper-level* counts (one paper contributes at most once to a given bin), (ii) *technique-mention* counts (normalized mentions aggregated across

papers), and (iii) *residual-observation* counts (distinct recoverable artifact instances reported by mobile studies). RQ1 uses both (i) and (ii), RQ2 uses (i), and RQ3 uses (iii).

A total of 197 primary studies published between 2005 and 2024 met the inclusion criteria. Figure 2 shows that AF research emerged slowly before 2010, gained traction between 2010 and 2013, and then increased steadily from 2014 onward, peaking in 2019 with 23 studies. The subsequent dip to 17 studies in 2020 likely reflects disruptions to academic workflows during the COVID-19 pandemic, with output rebounding to 21 studies in 2021. From 2020 through 2024, outputs remained relatively stable (13–21 studies/year), suggesting that AF has transitioned from an emerging topic to a mature and sustained research area.

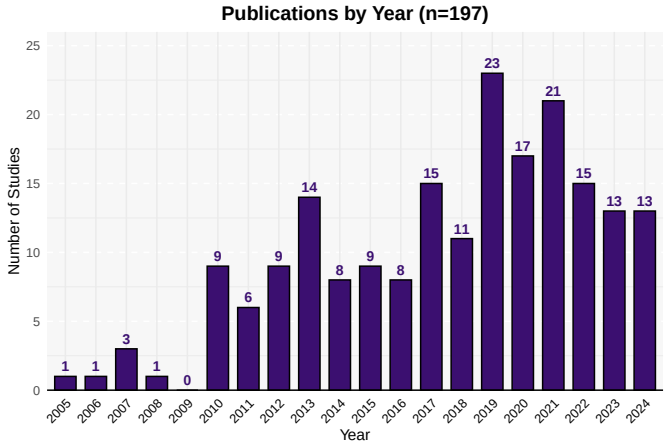

Figure 2: Publication of anti-forensics studies published by year.

5.1. RQ1: AF Techniques Examined in the Literature.

This subsection answers RQ1 by using our taxonomy to (i) characterize the *strategy families* of AF (high-level categories) and their evolution, (ii) examine how these strategies materialize as *concrete operations* (sub-techniques), and (iii) contextualize technique choices by domain and venue. We analyze AF using the strategy families and sub-techniques defined in Section 4. Across the 197 papers, we extracted 938 technique mentions and mapped each mention to exactly one high-level category and one normalized sub-technique (normalization rules in Section 3).

5.1.1. Trends in High-Level Categories.

Figure 3 summarizes how the *taxonomy families* appear over time. For each year and category, we count unique papers that report at least one technique in that category (a paper contributes at most once per category per year). Three families dominate across the entire period: Trail Obfuscation (TO) (118), Data Hiding (DH) (72), and Artifact Wiping (AW) (68). In contrast, Attacks Against Forensic Tools and Methods (AT) appears intermittently (39). Finally, PIA (21) and AAF (18) are largely post-2019 phenomena in our corpus.

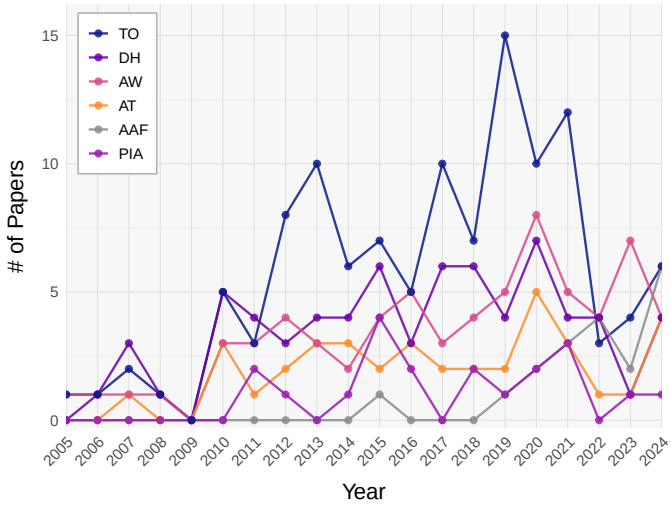

Figure 3: Yearly counts for six categories: Artifact Wiping (AW), Trail Obfuscation (TO), Data Hiding (DH), Possible Indications of Anti-Digital Forensics (PIA), Attacks Against Forensic Tools and Methods (AT), AI-Driven Anti-Forensics (AAF).

Following the AAF definition and subtypes in Section 4, we identify 18 AAF studies in the corpus. Of the 18 AAF studies, 12 are labeled as *Evasion*, 5 as *Generation*, and 3 as *Trace suppression* (subtype labels are multi-label; two studies fall under both Evasion and Generation). None implements and evaluates AAF as *mobile-device* anti-forensics in our corpus.

5.1.2. Trends in Sub-Techniques

While high-level categories reflect stable *strategic intent*, sub-techniques reveal *how* the literature operationalizes those intents. Figure 4 shows the temporal distribution of the 15 most prevalent sub-techniques. The pre-2010 period is sparse and narrow. From 2010 onward, multiple sub-techniques persist across years rather than appearing as isolated one-offs.

Data fabrication is the most persistent and intensive sub-technique ($n = 146$ total mentions), with a pronounced surge from 2019 to 2021 (e.g., 32 mentions in 2019 and 21 in 2021). Across the top-15 sub-techniques, 90/146 (61.6%) of all Data fabrication mentions occur between 2019-2024. *File wiping* (57) remains modest for much of the timeline but rises sharply in the early 2020s, peaking at 12 mentions in 2023 (Joun et al., 2023; Abozaid et al., 2023). Similarly, *log manipulation* (49) concentrates in a mid-2010s burst (peak 10 in 2015).

Media-pipeline sub-techniques (e.g., *JPEG anti-forensics* and *steganography*) are observed across multiple years (Stamm and Liu, 2010; Sharma et al., 2020; Chen et al., 2020). Finally, AI-oriented *adversarial example evasion* emerges only after 2019 and then recurs into 2024 (Wu et al., 2020; Li et al., 2021). In the counts, *adversarial example evasion* appears only after 2019 (21 total mentions), with peaks of 6 mentions in 2021, 2022, and 2024.

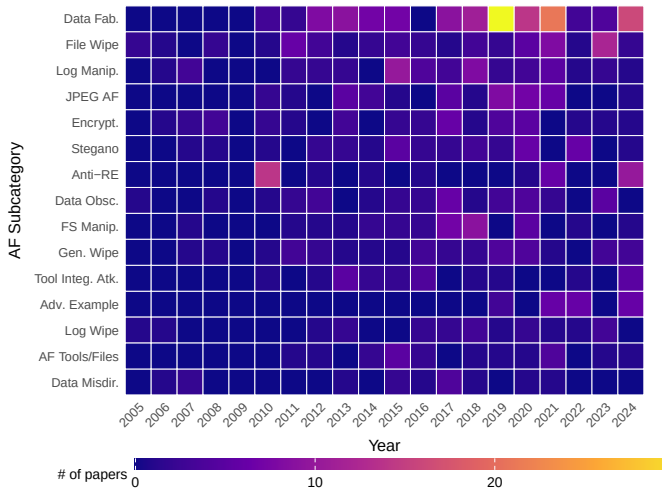

Figure 4: Heatmap of yearly mention counts for 15 anti-forensics subcategories. Abbrev.: Data fab. (Data fabrication); File wipe (File wiping); Log manip. (Log manipulation); JPEG AF (JPEG compression anti-forensics); Stegano (Steganography); Encrypt. (Encryption); Data obsc. (Data obscurity); Anti-RE (Anti-reverse engineering); FS manip. (Filesystem manipulation); Gen. wipe (Generic data wiping); Tool integ. atk. (Forensic software integrity attacks); AF tools/files (Presence of anti-forensics tools/files/installation files); Adv. example (Adversarial example, evasion); Log wipe (Log wiping); Data misdir. (Data misdirection/misinformation/framing).

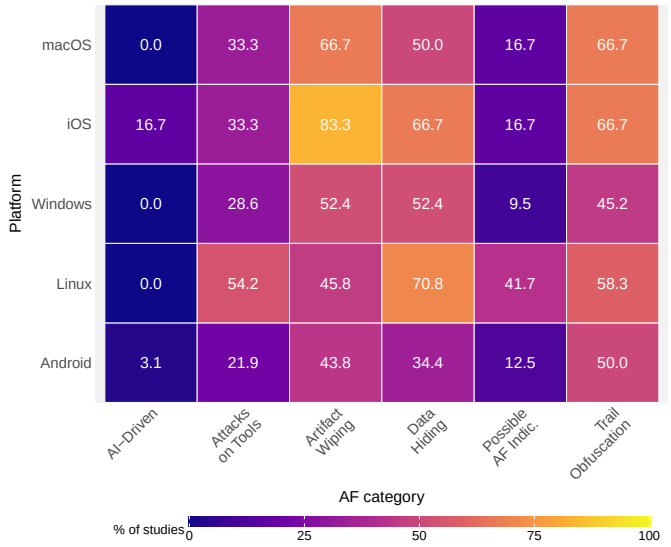

Figure 5: Heatmap of anti-forensics categories by platform (cell values are the percentage of studies on that platform using each category).

5.1.3. Domain and Publisher Distribution

To contextualize why certain taxonomy families dominate, we map each paper to a primary application domain (Appendix A, Table A.4) . *Multimedia (image/video/audio) (MM)* accounts for the largest share (83 papers), followed by *FS* (37) and *Mobile (MB)* (35). *General (GN)* totals 20; *Internet of Things (IoT)* totals 9; *Cloud (CL)* totals 2; and smaller slices appear in database (4), network (3), memory (3), and malware (1). Within MM, 27 studies use learning-based or generative methods in multimedia settings (24 image, 1 mobile, 2 cross-domain). Most studies appear in engineering-focused venues (IEEE and Elsevier together account for 146/197).

5.2. RQ2: How is AF research distributed across platforms, and to what extent do studies address mobile devices?

To evaluate platform focus, all 197 papers were categorized into four groups: *Mobile* (Android/iOS only), *Non-mobile* (desktop/server OS only), *Cross-platform* (explicitly compares mobile and non-mobile or multiple OS families), and *General* (no specific OS or device type; often surveys, conceptual work, or OS-agnostic multimedia studies). As shown in (Appendix B,Table B.5), only 31 studies (15.7%) examine mobile devices exclusively, while 46 (23.4%) focus on non-mobile environments. Cross-platform evaluations are rare (6 studies; 3.0%). The largest slice is General (114; 57.9%).

Beyond device type, we also examine which operating systems are explicitly studied. Platform labels are multi-label (a paper may cover multiple OSes): Windows is the most frequently studied (42), followed by Android (32) and Linux (24). Closed ecosystems (iOS and macOS, 6 each) are underrepresented.

5.2.1. Distribution of AF Categories Across Platforms

Figure 5 answers the taxonomy-specific part of RQ2: given a platform, which strategy families are most commonly studied? For each platform, we compute the percentage of platform-specific papers that mention at least one technique in each category (platform totals: Windows $n = 42$, Android $n = 32$, Linux $n = 24$, iOS $n = 6$, macOS $n = 6$). Because both platform and category labels are multi-label, a single paper may contribute to multiple categories for the same platform.

AW and *TO* are prevalent across all platforms (AW: 43.8% to 83.3%; TO: 45.2% to 66.7%). *DH* is especially common on Linux (70.8%) and iOS (66.7%). *AT* appears most frequently in Linux-focused work (54.2%) and less often in Android studies (21.9%). *PIA* remains relatively rare across platforms (9.5% to 41.7%), and AAF appears in only a small fraction of studies (Android 3.1%; iOS 16.7%).

5.2.2. Temporal Shifts in Mobile vs. Non-Mobile Research

(Appendix C, Figure C.7) shows the cumulative growth of studies by platform focus. Non-mobile and General studies accumulated steadily from 2005, while mobile-focused research remained nearly absent until 2013. From 2013 onward, mobile studies grew more consistently but at a slower rate than other categories. By 2024, mobile research ($n = 31$) had accumulated to approximately half the volume of non-mobile work ($n = 46$), with cross-platform studies ($n = 6$) remaining rare throughout. This pattern confirms that mobile AF research is a relatively recent and still-developing area compared to desktop-oriented investigations.

5.2.3. Cross-Platform Validation and Coverage of Techniques

Finally, we assess how often sub-techniques are validated across multiple operating systems. Some methods are broadly portable (e.g., *file wiping*, *generic data wiping*, *data obscurity*), while more specialized techniques remain platform-constrained. Overall, 43.9% of techniques are reported on three or more platforms, 31.6% on two, and 24.6% on only one, highlighting uneven cross-platform validation and limited generalizability.

5.3. RQ3: What forensic artifacts persist after AF activity on mobile devices, and how can they be categorized?

RQ3 uses the taxonomy to connect *anti-forensic strategy families* to *where evidence still persists*. To answer it, we filter the corpus to mobile studies with measurable artifact outcomes and then systematize residual evidence along two axes: (i) artifact *store type* (e.g., SQLite/WAL, logs, metadata, backup container) and (ii) *storage location bucket* (app-private, OS-level, external or shared, backup and cloud).

We first identify the mobile-only subset ($n = 31$) (Appendix B, Table B.5). From these, we retain papers that implemented and evaluated at least one AF technique on a mobile platform (excluding 4 mobile-only studies that did not evaluate an AF technique), yielding an RQ3-screened set of 27 studies. We then keep only studies that reported *recoverable residual artifacts* after the AF activity (excluding 9 studies without measurable residual outcomes), resulting in 18 studies used for the RQ3 results.

Across these 18 studies, we extract 93 residual artifact observations, where each observation represents one distinct recoverable residue instance explicitly reported by the primary study. (Appendix D, Table D.6) provides the complete per-study breakdown, detailing platform, AF category, artifact store type, storage location, and persistence characteristics for each of the 18 studies. Each observation is coded by platform (Android/iOS), artifact store type (SQLite/WAL, logs, caches and thumbnails, filesystem metadata, XML/plist, other, *backup containers*), location bucket (app-private, OS-level, external or shared, backup and cloud), and associated AF category.

5.3.1. Residual store types and storage locations

Residual observations most often persist in *app-private storage* (48; 51.6%), followed by *OS-level* locations (27; 29.0%). Notably, *backup and cloud* sources also appear (13; 14.0%), while *external or shared* storage is least represented in this subset (5; 5.4%).

Aggregating across locations, recoverable residues are dominated by the heterogeneous *Other* store type (32/93; 34.4%). In our coding, *Other* captures residual stores that do not map cleanly to SQLite/WAL, Logs, Cache and Thumbnails, FS-metadata, XML/plist, or backup containers, and primarily reflects media/content remnants and packaged containers (e.g., photos, videos, audio, ExoPlayer `.exo` fragments, encrypted/packaged files, vendor

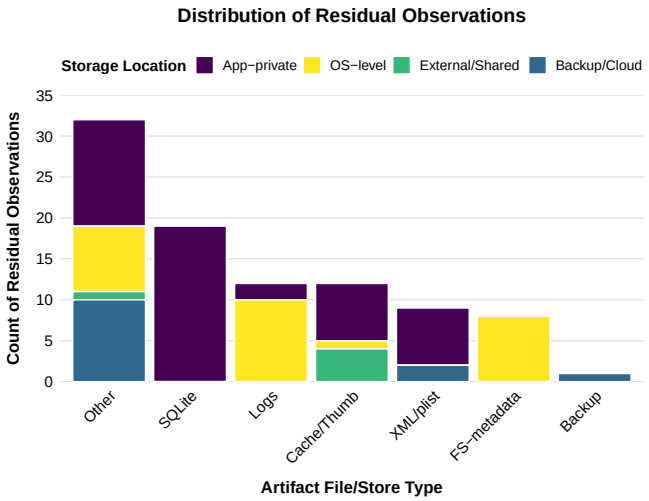

Figure 6: Distribution of residual observations by artifact store type and storage location (18-study residual-evidence subset; $n = 93$ observations).

backup containers such as `.edb/.senc/.ebk`, and tool artifacts) when explicitly reported.

SQLite/WAL accounts for 19/93 (20.4%) observations and appears exclusively in app-private storage in this subset. *Logs* (12/93; 12.9%) and *Cache and Thumbnails* (12/93; 12.9%) account for persistent telemetry and preview artifacts. *XML/plist* contributes 9/93 (9.7%), and *FS-metadata* contributes 8/93 (8.6%). Explicit *Backup containers* as a *store type* are rare (1/93; 1.1%). Figure 6 summarizes these distributions by store type across location buckets.

5.3.2. Mapping AF Categories to Residual Stores

Table 3 links AF *strategy families* to *where residue remains*. Entries are *paper-level co-occurrences*: a single study may contribute to multiple categories and multiple store types, so totals can exceed the number of unique residual observations. AW has the broadest co-occurrence footprint, appearing most frequently with *Other* stores and also recurring in *SQLite/WAL* and *Logs*. DH and TO co-occur across *Other*, *Cache and Thumbnails*, *FS-metadata*, and *SQLite/WAL*. AT appears infrequently in this subset, with occasional co-occurrences in *SQLite/WAL*, *Logs*, and *FS-metadata*. No PIA or AAF residues are observed in this mobile residual-evidence subset.

5.3.3. Platform-Level Distribution of Residues

Finally, we stratify residues by mobile platform. Android accounts for 84 of 93 residual observations across all store types, with the largest shares in *Other* (29) and *SQLite/WAL* (17), followed by *Logs* (12) and *Cache and Thumbnails* (10). *XML/plist* (8) and *FS-metadata* (7) contribute smaller but notable shares, while *Backup and Cloud* containers are rare (1). In contrast, iOS residues total only 9 observations, concentrated in *Other* (3), *SQLite/WAL* (2), and *Cache/Thumbnails* (2), with

Table 3: Mobile residual artifacts by store type × AF category.

File/Store Type	DH	AW	TO	AT	PIA	AAF	Row Σ
Cache/Thumbnails	3	2	1	0	0	0	6
FS-metadata	3	3	1	1	0	0	8
Logs	1	4	3	1	0	0	9
Other	4	8	5	0	0	0	17
SQLite/WAL	3	4	3	2	0	0	12
XML/plist	2	2	1	0	0	0	5
Backup containers	0	1	1	0	0	0	2
Column Σ	**16**	**24**	**15**	**4**	**0**	**0**	**59**

Legend: DH = Data Hiding; AW = Artifact Wiping; TO = Trail Obfuscation; AT = Attacks Against Forensic Tools and Methods; PIA = Possible Indications of Anti-Digital Forensics; AAF = AI-driven anti-forensics.

single observations in *FS-metadata* and *XML/plist*, and none in *Logs* or *Backup and Cloud*.

6. Discussion

Our longitudinal analysis reveals that AF has matured from a niche concern into a sustained research area. The publication timeline (Figure 2) shows three distinct phases: an emergent period prior to 2010 with sporadic contributions, a growth phase from 2010 to 2018 as the community established foundational techniques, and a mature phase following 2019 with stable output (13–23 studies/year). This stabilization implies that AF is no longer an emerging topic but an established subfield with ongoing research investment.

More significant than volume is the operational shift observed. While Trail Obfuscation ($n = 118$), Data Hiding ($n = 72$), and Artifact Wiping ($n = 68$) remain dominant (Figure 3), the methods adversaries use have evolved substantially. Data fabrication emerged as the most intensive sub-technique ($n = 146$ mentions), with 61.6% concentrated in 2019-2024 (Figure 4). This shift from *erasure* to *deception*, creating plausible false evidence rather than leaving suspicious gaps, reflects an increasingly sophisticated adversarial strategy. Classical operations such as file wiping and log manipulation continue to exist, but appear in episodic bursts, often coinciding with new forensic tools. Platform-level changes are also evolving this landscape; recent Android versions have deprecated APIs used by third-party wiping applications, potentially shifting artifact wiping toward vendor-controlled services.

The post-2019 emergence of AAF (18 studies) signals evolution toward AI-enabled anti-forensics. Among the three subtypes identified in Section 4, Evasion dominates (12 studies), using adversarial training and transferable black-box attacks to defeat forensic detectors (Ding et al., 2021; Wang et al., 2021; Li et al., 2021; Fang and Stamm, 2023). This dominance suggests that the current AAF threat model prioritizes evading automated detection over fabricating content, a pattern that may shift as on-device generative models become more accessible. Trace suppression (Liu et al., 2023; Wu and Sun, 2021) and generation-centric approaches (Peng et al., 2022; Chen et al., 2019) remain less explored but represent growing vectors as diffusion models and on-device AI capabilities mature. However, all 18 studies cluster exclusively in multimedia pipelines image manipulation, GAN-generated content, and compression artifact analysis (see Appendix A,Table A.4). Critically, none implements AI-driven anti-forensics on mobile platforms.

This gap is structural: mobile devices present different attack surfaces, sandboxed applications, proprietary filesystems such as Apple's APFS, and hardware-backed security that differ fundamentally from the desktop-centric multimedia pipelines where AAF has been validated.

KF1. AF research has evolved from classical erasure toward AI-driven deception, yet remains concentrated in multimedia domains, with no AAF validation on mobile platforms.

Research Gap (RG4). All 18 AAF studies focus on multimedia pipelines; none targets mobile platforms. Mobile-specific attack surfaces sandboxed apps, proprietary filesystems, hardware-backed security remain untested against AI-driven anti-forensics.

TThe platform distribution (Section 5, Table B.5) reveals a fundamental imbalance: mobile-only studies constitute only 15.7% of the corpus, compared to 23.4% for desktop and server environments, while the majority (57.9%) remain platform-agnostic surveys, taxonomies, or multimedia work that cannot be applied directly to device-specific acquisition scenarios. Cross-platform validation is particularly scarce, with only 3.0% of studies explicitly comparing mobile and non-mobile environments.

Operating system coverage compounds this imbalance (Figures 5) and (Appendix C C.7). Windows and Android dominate ($n = 42$ and $n = 32$ respectively), indicating both data accessibility and the relative ease of experimental testing on open platforms. Closed ecosystems remain severely underexplored: iOS and macOS each appear in only 6 studies, largely due to proprietary restrictions that limit experimental replication. This disparity creates a methodological blind spot forensic practitioners increasingly encounter iOS devices in casework, yet the evidence base for iOS-specific AF countermeasures remains thin.

The category-by-platform heatmap (Figure 5) distinguishes portable from platform-specific techniques. Artifact Wiping (43.8% to 83.3%) and Trail Obfuscation (45.2% to 66.7%) appear consistently across all platforms, suggesting that core AF behaviors erasure, concealment, and trace manipulation are replicated rather than redefined in mobile contexts. Data Hiding is more common on Linux (70.8%) and iOS (66.7%), which is indicative of filesystem-level concealment adapted to these operating systems. Attacks on Tools are more common in Linux-

focused work (54.2%) than on Android (21.9%) which is a notable gap given Android's dominance in mobile forensics.

KF2. Desktop and open platforms (Windows, Android) dominate AF research, while closed ecosystems (iOS, macOS) remain largely unexplored. Cross-platform validation is rare (3.0%).

The most striking finding from RQ3 is that mobile anti-forensics rarely achieves complete erasure. Across 18 studies, we extracted 93 residual artifact observations, evidence persisting despite deliberate AF activity (Figure 6, Table 3). Instead of destruction, evidence is displaced into secondary and structured containers SQLite databases, WAL files, system logs, and filesystem metadata, this provides important opportunities for partial reconstruction of user actions and AF behavior.

Residual evidence concentrates in predictable locations. 51.6% of observations are related to app-private storage, which shows the forensic value of application sandboxes where SQLite databases and WAL files persist after deletion. OS-level locations (29.0%) capture system logs and filesystem metadata that applications cannot delete. As mobile backup ecosystems grow, backup and cloud artifacts (14.0%) represent an understudied recovery path.

Store-type distribution guides recovery priorities. SQLite/WAL (20.4%) remains a high-value target write-ahead logs retain deleted records that primary database queries miss (Salamh et al., 2020; Heath et al., 2023). Logs and Cache/Thumbnails (12.9% each) preserve notification metadata and media previews surviving content deletion. The heterogeneous "Other" category (34.4%) captures diverse media remnants and vendor-specific containers that vary across devices and lack unified documentation.

The category-to-store mapping (Table 3) yields concrete insight. Artifact Wiping leaves the largest residual footprint, which is found in "Other" and SQLite/WAL confirming wiping tools target primary payloads while secondary containers escape. Data Hiding and Trail Obfuscation frequently generate recoverable traces in caches, logs, and metadata, indicating that deliberate concealment or obfuscation fails to remove secondary data structures created by the OS or app frameworks.

Platform stratification reveals a critical gap. Android dominates residual observations across all store types; iOS residues are few (9 total), with no log-based traces documented. Due to this imbalance, which is a result of both corpus bias and iOS acquisition limitations, residual patterns unique to iOS are not well described.

KF3. Mobile AF displaces rather than erases evidence. Residual traces persist in SQLite/WAL, logs, and filesystem metadata especially, on Android. Recognizing where these residuals persist is critical for improving mobile triage and counter-AF validation.

In light of these findings, investigators should prioritize sandbox-aware acquisition with explicit WAL parsing and log extraction on Android, and concentrate on metadata and preview artifacts on iOS, where traces are more likely to persist.

7. Limitation

This SoK has several limitations. We used Scopus as the primary database; relevant grey literature or non-indexed studies may have been missed. Technique classification was primarily performed by one researcher, and coding decisions were reviewed during analysis by co-authors, but formal inter-rater reliability was not assessed. The corpus reflects existing research emphases rather than real-world AF prevalence. RQ3 findings derive from only 18 studies (93 observations), heavily Android-biased, limiting generalizability of residual evidence patterns to iOS. Finally, we report what primary studies documented; we did not independently validate AF tool effectiveness or artifact persistence.

8. Conclusion

This SoK analyzed 197 studies to systematize anti-forensics techniques, platform coverage, and mobile residual evidence. Three principal findings emerge. First, AF has evolved from traditional erasure toward deception-oriented and AI-enabled approaches. While Data Hiding, Artifact Wiping, and Trail Obfuscation remain dominant, data fabrication now exceeds file wiping as the primary operational technique, and AAF has emerged as a distinct category though only used in multimedia pipelines, with no mobile-platform validation. Second, platform coverage is fundamentally imbalanced. Windows and Android dominate ($n = 42$ and $n = 32$), while iOS and macOS each appear in only 6 studies. Cross-platform validation remains rare (3.0%), which limits countermeasure generalizability. Third, mobile AF displaces rather than erases evidence. Across 18 studies reporting residual outcomes, 93 artifact observations persist in SQLite/WAL, logs, caches, and filesystem metadata predominantly in app-private (51.6%) and OS-level (29.0%) storage.

These findings yield three recommendations: (i) develop standardized cross-platform benchmarks to enable reproducible evaluation; (ii) expand AAF research to mobile platforms, where on-device AI capabilities are proliferating; and (iii) prioritize iOS-focused studies to address the current evidence gap. For practitioners, the residual evidence patterns identified here provide concrete guidance: sandbox-aware acquisition with explicit WAL parsing on Android, and metadata-focused recovery on iOS. This SoK provides an empirical baseline for AF research and identifies where evidence persists despite adversarial activity, offering a foundation for more resilient forensic practice.

References

Fatemah Abozaid, Sharifa Yousif, M Abdulhamid Shafi'i, Dafalla Izaruku Zubeir, and Abdullah Hussein Al-Ghushami. Deployment and detection of anti-forensics shredding technique in android mobile devices. In *2023 Sixth International Conference of Women in Data Science at Prince Sultan University (WiDS PSU)*, pages 150–155. IEEE, 2023.

Pietro Albano, Aniello Castiglione, Giuseppe Cattaneo, Giancarlo De Maio, and Alfredo De Santis. On the construction of a false digital alibi on the android os. In *2011 Third International Conference on Intelligent Networking and Collaborative Systems*, pages 685–690. IEEE, 2011a.

Pietro Albano, Aniello Castiglione, Giuseppe Cattaneo, and Alfredo De Santis. A novel anti-forensics technique for the android os. In *2011 International Conference on Broadband and Wireless Computing, Communication and Applications*, pages 380–385. IEEE, 2011b.

Fahdiaz Alief, Yohan Suryanto, Linda Rosselina, and Tofan Hermawan. Analysis of autopsy mobile forensic tools against unsent messages on whatsapp messaging application. In *2020 7th International Conference on Electrical Engineering, Computer Sciences and Informatics (EECSI)*, pages 26–30. IEEE, 2020.

Cosimo Anglano. Forensic analysis of whatsapp messenger on android smartphones. *Digital Investigation*, 11(3): 201–213, 2014.

Shiva Azadegan, Wei Yu, Hui Liu, M Sistani, and Subrata Acharya. Novel anti-forensics approaches for smart phones. In *2012 45th Hawaii International Conference on System Sciences*, pages 5424–5431. IEEE, 2012.

Konstantia Barmpatsalou, Dimitrios Damopoulos, Georgios Kambourakis, and Vasilios Katos. A critical review of 7 years of mobile device forensics. *Digital Investigation*, 10(4):323–349, 2013.

Konstantia Barmpatsalou, Tiago Cruz, Edmundo Monteiro, and Paulo Simoes. Current and future trends in mobile device forensics: A survey. *ACM Computing Surveys (CSUR)*, 51(3):1–31, 2018.

H Hemanth Bharath Bhushan and S Metilda Florance. An overview on handling anti forensic issues in android devices using forensic automator tool. In *2022 IEEE International Conference on Signal Processing, Informatics, Communication and Energy Systems (SPICES)*, volume 1, pages 425–430. IEEE, 2022.

Chang Chen, Zhiwei Xiong, Xiaoming Liu, and Feng Wu. Camera trace erasing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2950–2959, 2020.

Chen Chen, Xinwei Zhao, and Matthew C Stamm. Generative adversarial attacks against deep-learning-based camera model identification. *IEEE Transactions on Information Forensics and Security*, 2019.

Kevin Conlan, Ibrahim Baggili, and Frank Breitinger. Anti-forensics: Furthering digital forensic science through a new extended, granular taxonomy. *Digital investigation*, 18:S66–S75, 2016.

Feng Ding, Guopu Zhu, Yingcan Li, Xinpeng Zhang, Pradeep K Atrey, and Siwei Lyu. Anti-forensics for face swapping videos via adversarial training. *IEEE Transactions on Multimedia*, 24:3429–3441, 2021.

Alessandro Distefano, Gianluigi Me, and Francesco Pace. Android anti-forensics through a local paradigm. *digital investigation*, 7:S83–S94, 2010.

Gokila Dorai, Sudhir Aggarwal, Neet Patel, and Charisa Powell. Vide-vault app identification and extraction system for ios devices. *Forensic Science International: Digital Investigation*, 33:301007, 2020.

Christian DOrazio, Aswami Ariffin, and Kim-Kwang Raymond Choo. ios anti-forensics: How can we securely conceal, delete and insert data? In *2014 47th Hawaii International Conference on System Sciences*, pages 4838–4847. IEEE, 2014.

Shengbang Fang and Matthew C Stamm. Attacking image splicing detection and localization algorithms using synthetic traces. *IEEE transactions on information forensics and security*, 19:2143–2156, 2023.

Aya Fukami, Sasha Sheremetov, Francesco Regazzoni, Zeno Geradts, and Cees De Laat. Experimental evaluation of e. mmc data recovery. *IEEE Transactions on Information Forensics and Security*, 17:2074–2083, 2022.

Thomas Göbel and Harald Baier. Anti-forensics in ext4: On secrecy and usability of timestamp-based data hiding. *Digital Investigation*, 24:S111–S120, 2018.

Murat Gül and Emin Kugu. A survey on anti-forensics techniques. In *2017 International Artificial Intelligence and Data Processing Symposium (IDAP)*, pages 1–6. IEEE, 2017.

Ryan Harris. Arriving at an anti-forensics consensus: Examining how to define and control the anti-forensics problem. *digital investigation*, 3:44–49, 2006.

Saeed Shafiee Hasanabadi, Arash Habibi Lashkari, and Ali A Ghorbani. A survey and research challenges of anti-forensics: Evaluation of game-theoretic models in simulation of forensic agents' behaviour. *Forensic Science International: Digital Investigation*, 35:301024, 2020.

Howard Heath, Áine MacDermott, and Alex Akinbi. Forensic analysis of ephemeral messaging applications: Disappearing messages or evidential data? *Forensic Science International: Digital Investigation*, 46:301585, 2023.

Jihun Joun, Sangjin Lee, and Jungheum Park. Discovering spoliation of evidence through identifying traces on deleted files in macos. *Forensic Science International: Digital Investigation*, 44:301502, 2023.

Karl-Johan Karlsson and William Bradley Glisson. Android anti-forensics: Modifying cyanogenmod. In *2014 47th Hawaii International Conference on System Sciences*, pages 4828–4837. IEEE, 2014.

Gary C. Kessler. Anti-forensics and the digital investigator. In *Proceedings of the Australian Digital Forensics Conference*, Perth, Australia, 2007. Edith Cowan University.

Jiyoun Kim, Seunghee Seo, Byoungjin Seok, and Changhoon Lee. Analysis of digital forensics artifacts based on deletion methods in mobile wiping applications. In *2024 International Conference on Platform Technology and Service (PlatCon)*, pages 177–181. IEEE, 2024.

Barbara Kitchenham et al. Procedures for performing systematic reviews. *Keele, UK, Keele University*, 33(2004): 1–26, 2004.

Dongze Li, Wei Wang, Hongxing Fan, and Jing Dong. Exploring adversarial fake images on face manifold. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5789–5798, 2021.

Chi Liu, Huajie Chen, Tianqing Zhu, Jun Zhang, and Wanlei Zhou. Making deepfakes more spurious: evading deep face forgery detection via trace removal attack. *IEEE Transactions on Dependable and Secure Computing*, 20(6):5182–5196, 2023.

Nemanja Maček, Perica Štrbac, Dušan Čoko, Igor Franc, and Mitko Bogdanoski. Android forensic and anti-forensic techniques–a survey. In *8th International Conference on Business Information Security (BISEC'2016)*, 2016.

Walter T Mambodza and AR Nagoormeeran. Android mobile forensic analyzer for stegno data. In *2015 International Conference on Circuits, Power and Computing Technologies [ICCPCT-2015]*, pages 1–8. IEEE, 2015.

Mohammad M Mirza, Fahad E Salamh, and Umit Karabiyik. An android case study on technical anti-forensic challenges of whatsapp application. In *2020 8th International Symposium on Digital Forensics and Security (ISDFS)*, pages 1–6. IEEE, 2020.

Julián García Murias, Douglas Levick, and Sean McKeown. A forensic analysis of streaming platforms on android os. *Forensic Science International: Digital Investigation*, 44:301485, 2023.

Christopher Neale. Fool me once: A systematic review of techniques to authenticate digital artefacts. *Forensic Science International: Digital Investigation*, 45:301516, 2023.

Myungseo Park, Hangi Kim, and Jongsung Kim. How to decrypt pin-based encrypted backup data of samsung smartphones. *Digital Investigation*, 26:63–71, 2018.

Fei Peng, Liping Yin, and Min Long. Bdc-gan: Bidirectional conversion between computer-generated and natural facial images for anti-forensics. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(10): 6657–6670, 2022.

Sebastián Potockỳ, Jozef Štulrajter, and Michal Turčaník. Advanced industrial espionage: Mobile device as anti-forensic and anonymization tool. In *2024 New Trends in Signal Processing (NTSP)*, pages 1–8. IEEE, 2024.

Marcus Rogers. Anti-forensics: the coming wave in digital forensics. *Retrieved September*, 7(2008):3, 2006.

Hao Ruan, Xiao Fu, Xuanyu Liu, Xiaojiang Du, and Bin Luo. Analyzing android application in real-time at kernel level. In *2017 26th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–9. IEEE, 2017.

Fahad E Salamh, Umit Karabiyik, and Marcus K Rogers. Asynchronous forensic investigative approach to recover deleted data from instant messaging applications. In *2020 International Symposium on Networks, Computers and Communications (ISNCC)*, pages 1–6. IEEE, 2020.

Shishir Sharma, Hareesh Ravi, A Venkata Subramanyam, and Sabu Emmanuel. Anti-forensics of median filtering and contrast enhancement. *Journal of Visual Communication and Image Representation*, 66:102682, 2020.

Matthew C Stamm and KJ Ray Liu. Wavelet-based image compression anti-forensics. In *2010 IEEE international conference on image processing*, pages 1737–1740. IEEE, 2010.

Vrizlynn LL Thing, Kian-Yong Ng, and Ee-Chien Chang. Live memory forensics of mobile phones. *digital investigation*, 7:S74–S82, 2010.

Yongwei Wang, Xin Ding, Yixin Yang, Li Ding, Rabab Ward, and Z Jane Wang. Perception matters: Exploring imperceptible and transferable anti-forensics for gan-generated fake face imagery detection. *Pattern Recognition Letters*, 146:15–22, 2021.

Jianyuan Wu and Wei Sun. Towards multi-operation image anti-forensics with generative adversarial networks. *Computers & Security*, 100:102083, 2021.

Jianyuan Wu, Li Liu, Xiangui Kang, and Wei Sun. A generative adversarial network framework for jpeg anti-forensics. In *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1442–1447. IEEE, 2020.

Xiaosong Zhang, Yu-an Tan, Yuan Xue, Quanxin Zhang, Yuanzhang Li, Can Zhang, and Jun Zheng. Cryptographic key protection against frost for mobile devices. *Cluster Computing*, 20(3):2393–2402, 2017.

Appendix A. Domain distribution across platform

Table A.4: Domain distribution across publishers with coefficient of variation (CV).

Publishers	Papers	MM	FS	MB	GN	IoT	DB	NW	MEM	CL	MW
IEEE	106	45	13	22	10	6	3	3	2	2	0
Elsevier	40	10	13	9	6	1	0	0	1	0	0
ACM	18	11	4	1	0	1	1	0	0	0	0
Springer	17	14	0	1	1	1	0	0	0	0	0
Others	16	3	7	2	3	0	0	0	0	0	1
Total	197	83	37	35	20	9	4	3	3	2	1

Legend: DH = Data Hiding; AW = Artifact Wiping; TO = Trail Obfuscation; AT = Attacks Against Forensic Tools and Methods; PIA = Possible Indications of Anti-Digital Forensics; AAF = AI-driven Anti-Forensics.

Appendix B. Device Type Distribution

Table B.5: Device type distribution across included studies

Device Type	Count	Percentage (%)
Mobile	31	15.7
Non-mobile	46	23.4
Cross-platform	6	3.0
General	114	57.9
Total	**197**	**100.0**

Appendix C. Cumulative fraction of studies by platform focus (Mobile, Non-mobile, Cross-platform, General).

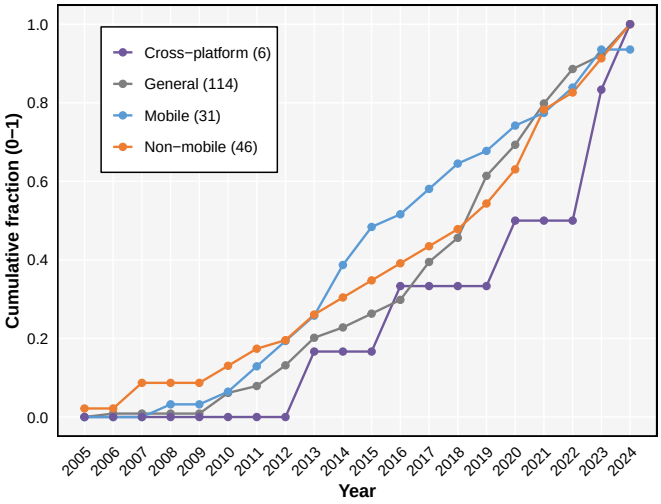

Figure C.7: Cumulative Fraction of studies by Platform Focus.

Appendix D. Mobile anti-forensic Residual Artifacts per study

Table D.6: Mobile anti-forensics residual artifacts per study (RQ3). One row per paper.

	Platform		*Technique Category*					*Artifact File/Store Type*						*Location*	*Persistence*	*Evidence (Extraction note)*
Paper (Year)	**Android**	**iOS**	Data Hiding	Artifact Wiping	Trail Obfusc.	Attack on Tools	Possible Indic.	**SQLite**	**XML/ plist**	**Logs**	Cache/ Thumb	FS-meta	**Other**			
Distefano et al. (2010)	✓	✗	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗	✓	APP	NR	Observed private-folder storage and DB size changes; uninstall removed artifacts.
Albano et al. (2011a)	✓	✗	✗	✓	✓	✗	✗	✗	✗	✓	✗	✗	✓	APP, EXT, OS	UR, VOL	Observed install/uninstall traces; observed SD data persists; logs volatile
Azadegan et al. (2012)	✓	✗	✗	✓	✓	✗	✗	✗	✗	✓	✗	✓	✗	OS	NR	Observed Logcat provider access and tool activity; persistence not tested
DOrazio et al. (2014)	✗	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓	✗	OS	RP	Observed default icons after reboot; journal unchanged; metadata edits identified.
Karlsson and Glisson (2014)	✓	✗	✗	✗	✗	✓	✗	✓	✗	✓	✗	✗	✗	APP, OS	NR	Observed provider system logs: caller and query arguments; Identified alternate contacts SQLite DBs for Cellebrite/XRY
Mambodza and Nagoormeeran (2015)	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	NR	NR	Observed stego scan results; persistence not tested
Zhang et al. (2017)	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	VOL, RP	NR	Compared pre/post-reboot RAM dumps; observed key overwritten, remanence retained
Ruan et al. (2017)	✓	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✓	✗	OS	NR	E1: Parsed kernel-level syscall logs to reconstruct behavior graph
Göbel and Baier (2018)	✓	✗	✓	✗	✗	✓	✗	✗	✗	✗	✗	✓	✗	OS	RP	Analyzed ext4 inode timestamps; observed delete leaves fields unchanged
Park et al. (2018)	✓	✗	✓	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓	Bkp	BKP	Identified PIN-based Manifest.exml; observed backup stored on PC
Alief et al. (2020)	✓	✗	✗	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	APP	NR	Observed WhatsApp traces in extracted database; imaged one week later
Mirza et al. (2020)	✓	✗	✗	✓	✗	✗	✗	✓	✗	✓	✗	✗	✗	APP, OS	NR	Observed capture via notification listener; stored recovered messages in app database
Dorai et al. (2020)	✗	✓	✓	✗	✗	✗	✗	✓	✗	✗	✓	✗	✓	APP	NR	Extracted app databases, recovered media and plists; persistence untested
Salamh et al. (2020)	✓	✗	✓	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗	APP	RP	Recovered WhatsApp WAL data; observed WAL persists after shutdown.
Fukami et al. (2022)	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	OS	RP	Observed raw-flash recovery unchanged after repeating power cycles
Heath et al. (2023)	✓	✓	✗	✓	✓	✗	✗	✗	✗	✗	✓	✗	✓	APP	RP	Retrieved deleted rows from db-wal; parsed msgstore.db timers
Murias et al. (2023)	✓	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	APP, OS	NR	Compared post-cache-clear recovery; matched deletion timestamps; attributed data loss to TRIM
Potockỳ et al. (2024)	✓	✗	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	OS	NR	Observed dumpsys audio and camera connect/disconnect service events

Legend: ✓present, ✗absent. **Location** codes: App = App-private, OS = OS-level, Ext = External/Shared, Bkp = Backup/Cloud.
Persistence. codes: RP = reboot-persistent, UR = uninstall-resistant, VOL = volatile, NR = Not reported RU = Removed on Uninstall.