CaseLinker: A Framework for Retrospective Analysis and Case Studies of Internet Crimes Against Children Across U.S. Task Forces, 2010–2026

Mrinaal Ramachandran
Graduate Student, Department of Computer Science, University of Massachusetts Amherst
Independent Research | github.com/mrinaalr/CaseLinker | Live Demo
Technical Report Series #4 | April 2026

Abstract

The history of the internet has been predominantly championed from the perspective of platforms, growth, and innovation. Largely absent from that record is a systematic account of how the same infrastructure transformed the landscape of child exploitation—and how law enforcement institutions adapted, strained, and responded across decades of technological change. This gap is not incidental: internet-facilitated crimes against children have scaled alongside the web itself, yet their history remains fragmented across jurisdictions, documented in press releases but never systematically connected at scale.

This paper presents a framework for the first longitudinal case study analysis of the Internet Crimes Against Children (ICAC) Task Force Program across 2010–2026, organized across four technological eras: the social platform era (2010–2014), the mobile-first transition (2015–2018), platform proliferation and expansion (2019–2022), and the social media and generative AI period (2023–2026). Drawing on 2,500 publicly available case reports processed through CaseLinker, we develop a structured methodology for 20 case studies that trace how perpetrator methodolo-

gies, platform surfaces, and law enforcement responses co-evolved across each era. Each case study is analyzed across five dimensions: platform context, perpetrator methodology, investigative approach, prosecutorial outcomes, and relationship to the broader technological moment. The framework, dataset, and analysis pipeline are released as open-source software designed to be replicable and extensible by investigators, researchers, and organizations working in this domain.

Keywords: ICAC investigations, retrospective case analysis, longitudinal study, internet history, child exploitation, open-source research infrastructure

---

Contents

1. Introduction

1.1. The Record That Does Not Exist

The internet turned thirty years old as a mass-market technology in 2024. In three decades, it moved from a research curiosity to the primary infrastructure of social, civic, and economic life for billions of people. That transformation has been extensively documented: the platforms, the companies, the cultural shifts, the policy battles.

What has not been documented—not systematically, not at scale—are online crimes and how bad actors have exploited this revolution to victimize children.

Internet-facilitated crimes against children did not emerge after the web matured. They emerged with it. The same compressed timeline that brought social platforms, mobile connectivity, and generative AI to billions of users brought those capabilities to perpetrators as well. Each wave of technological adoption created new vectors: early chat rooms and P2P networks gave way to social media, then mobile-first platforms, then encrypted messaging, then AI-generated content. Law enforcement adapted. Task forces were built, investigative methodologies evolved, prosecutorial outcomes accumulated. But no system has connected those outcomes across jurisdictions, eras, and agencies into a coherent historical record.

That gap is what this paper addresses.

1.2. The CaseLinker Series

This is the fourth report in the CaseLinker technical series. Report #1 introduced the system architecture: a modular five-layer pipeline for ingesting, processing, storing, cluster-

ing, and visualizing publicly available ICAC case reports, evaluated on 47 cases from Arizona ICAC annual reports [14]. Report #2 scaled to 207 cases across additional sources, introduced named entity recognition via Stanford Stanza for agency extraction, and documented the distributed law enforcement network structure visible at that scale [15]. Report #3 scaled further to 500 cases across five sources, introduced nine-dimensional facet tree navigation, and formalized the adversarial risk argument through a utility asymmetry proof [16].

This report does not introduce a new algorithmic component. Its contribution is different: it establishes the framework through which CaseLinker's corpus will be used to conduct structured historical case studies—the methodology, the sampling strategy, the analytical dimensions, the ethical and legal grounding, and the path to execution.

CaseLinker Series Snapshot (April 2026)

Reports published: 4
Cases processed: 2,500+
Jurisdictions: 25 U.S. task forces and 1000+ agencies
Features extracted: 27,150; 10 dimensions
HRPO Determination: #7668 (Not Human Subjects Research, 45 CFR 46)

1.3. Why Case Studies, Why Now

Three reports in, the argument for retrospective case analysis has been made and validated. Patterns emerge at scale that are invisible in individual cases: the distributed-core structure of ICAC agency networks, perpetrator age distributions, offense type composition, jurisdiction-level variation in investigation types. The analytical value compounds.

But aggregate statistics, however useful, do not capture the texture of individual investigations. The platform context matters. So does the perpetrator methodology, the relationship to the victim, the demographic reality of who was harmed — and the investigative approach: the triage decisions made under pressure, the agency coordination that turns a tip into a prosecution, the operational judgment that rescues children. The prosecutorial outcome, and what it reveals about the moment in which the case occurred, is the kind of evidence that moves policy — the concrete, human record that aggregate statistics cannot deliver to a legislator's desk. Statistics tell you about perpetrator ages and common platforms. Case studies tell you what that looked like in practice — which

platforms were involved, how a perpetrator made contact, how law enforcement built a case, and what it took to get a conviction. That is the texture of this crime. That is what aggregate numbers cannot hold.

The generative AI period creates additional urgency. The shift from 2022 to 2026 represents the most significant discontinuity in the technological landscape of child exploitation since the emergence of smartphones. AI-generated CSAM increased 1,325% between 2023 and 2024 alone [1]. Online enticement reports increased 192% in 2024 [1]. And this is likely the floor, not the ceiling — as generative tools become cheaper, more accessible, and harder to detect, the exploitation methodologies of the next five years will bear little resemblance to those of the last five. Understanding where this crime is going requires understanding where it has been — and that means systematically examining the record of successful enforcement before the discontinuity widens further. The platforms, methodologies, and perpetrator profiles associated with this crime are not yet well-characterized in the public record—precisely because the tools to study them at scale did not previously exist.

CaseLinker now has 2,500 cases. The case studies can begin.

2. Historical Framework: Four Technological Eras

The four eras used to organize this analysis are not arbitrary periodizations. Each represents a distinct technological shift that materially changed both the landscape of exploitation and the law enforcement response to it. The eras are defined by platform infrastructure, not calendar boundaries—the transition dates are approximate and cases near the boundaries may share characteristics of adjacent periods.

2.1. Era 1: The Social Platform Era (2010–2014)

The early 2010s represent the first period in which social media platforms achieved mass adoption among adolescents. Facebook's user base exceeded 1 billion in 2012. Instagram launched in 2010 and was acquired by Facebook in 2012. Snapchat launched in 2011 and introduced ephemeral content as a feature. These platforms transformed the social environment for young people in ways that outpaced both parental awareness and legal frameworks.

For ICAC investigations, this era is characterized by the emergence of platform-mediated contact as a primary grooming vector. The AZICAC dataset—CaseLinker's original source, covering 2011–2014—reflects this transition: "online" and "chat" appear as the dominant platform mentions in ten and seven cases respectively, with Facebook appearing in two cases in 2012 and 2014. Investigation types from this era vary from reactive operations triggered by parental reports and platform referrals, to 12 proactive and 3 undercover successful operations documented, demonstrating steady improvements as task forces developed methodologies for platform-based investigations.

2.2. Era 2: The Mobile-First Transition (2015–2018)

Smartphone prevalence among U.S. adolescents crossed 50% around 2015 and continued accelerating through 2018. This transition moved internet access from a shared family device to a personal, pocket-carried, always-on connection. The implications for exploitation were significant: the geographic and temporal constraints that had previously bounded predatory contact—a child had to be at a computer, at home, during limited hours—largely disappeared.

This era also saw the rise of encrypted and disappearing-message platforms as mainstream tools. Kik, Whisper, and early versions of disappearing-story features created new surfaces for contact initiation that were harder to detect and document. ICAC investigations from this period reflect a shift toward more complex digital forensics and challenging cases: multiple devices, multiple platforms, and increasingly sophisticated evidence chains.

2.3. Era 3: Platform Proliferation and Expansion (2019–2022)

The period from 2019 to 2022 is defined by the fragmentation of the platform landscape and the acceleration of volume. TikTok reached Western mass adoption after 2018. Discord—originally a gaming communication tool—became a general-purpose platform heavily used by adolescents. Livestreaming platforms created real-time contact surfaces. The number of distinct platforms appearing in ICAC investigations expanded substantially, with the facet search tree showing increased diversity in social media platforms in successful operations and 795 unique cases, a substantial increase from 172 and 94 in the previous eras respectively.

This era also saw COVID-19 dramatically accelerate online engagement among children and adolescents. School-from-home, social isolation, and increased screen time created conditions that exploitation research has documented as correlated with increased online

risk. CyberTipline report volume grew substantially across this period, reflecting both increased exploitation activity and improved platform detection infrastructure.

2.4. Era 4: Social Media and Generative AI (2023–2026)

The current era is characterized by two convergent developments: the maturation of short-form social media as the dominant adolescent platform environment, and the emergence of generative AI as both a threat vector and a detection challenge. NCMEC documented 67,000 AI-CSAM CyberTip reports in 2024, up from 4,700 the prior year—a 1,325% increase in a single year [1]. Financial sextortion emerged as a distinct and rapidly scaling crime typology, with 546,000 online enticement reports in 2024 alone. There are 1147 reports with operations completed in the past 3 years.

Law enforcement response in this era reflects both the accumulated capability of 25 years of ICAC infrastructure and the genuine strain of volume. The REPORT Act of 2024 expanded mandatory reporting obligations. Platform cooperation has become more variable as E2EE deployment creates detection blind spots. The investigative challenges of an ever evolving digital world continue to compound—with new exploitation surfaces and complex investigations—representing the dynamic challenges of protecting children and holding abusers accountable in today's computerized world.

Table 1: Four-Era Framework Summary

Era	**Period**	**Defining Technology**	**LE Adaptation**
Social Platform	2010–2014	Facebook, Instagram, Snapchat reach mass adolescent adoption	Platform-mediated contact becomes primary grooming vector; reactive investigations dominant
Mobile-First	2015–2018	Smartphone prevalence >50%; encrypted/disappearing messaging	Device forensics complexity increases; undercover methodology expands
Platform Proliferation	2019–2022	TikTok, Discord, livestreaming; COVID-accelerated online engagement	Investigation volume surges; multi-platform evidence chains become standard
Generative AI	2023–2026	AI-generated CSAM; financial sextortion; E2EE expansion	Hash-matching blind spots; REPORT Act; new crime typologies

3. Research Methodology

3.1. Grounded in CaseLinker

This research is grounded in CaseLinker, an open-source retrospective case analysis system developed to process, structure, and surface patterns from publicly available ICAC press releases and task force case reports. The system is auditable, interpretable, and appropriate for the sensitivity of the subject matter—operating exclusively on already-public, already-redacted case summaries and producing no outputs that identify individuals. The full source corpus, extraction code, and analysis pipeline are released as open-source software to support replication and extension by other researchers.

At the time of this report, CaseLinker has processed 2,500+ publicly available case reports across 25 U.S. jurisdictions, extracting features across ten dimensions: topic, severity, platform, investigation type, source, agency, prosecution outcome, location, severity phrase, and date-range. The 20 case studies in this series will be selected from this corpus using the selection criteria described below.

3.2. Case Study Selection: Stratified Sampling

The 20 case studies will be selected using stratified sampling across the four historical eras and five analytic dimensions. The goal is not statistical generalization from the sample, but grounded historical illustration of patterns and institutional responses that aggregate statistics and individual case reading cannot reveal in combination.

Era distribution: Cases will be sampled across all four eras in approximate proportion to the corpus coverage of each period, with a minimum of four cases per era to ensure each technological moment is represented.

Analytic dimensions: Within each era, cases are selected to represent variation across:

1. **Platform involvement** — from early social media contact to encrypted messaging to recent online sextortion
2. **Victim and perpetrator demographics** — age, relationship, relevant context
3. **Investigation type** — proactive, reactive, and undercover operations
4. **Geography and jurisdiction** — variation across task forces, states, and agency

networks

5. **Agency coordination** — single-agency cases versus multi-agency and large scale operations

Cases are identified using CaseLinker's ten-dimensional facet tree, which enables precise filtering by era, platform, investigation type, severity, and source. The facet tree converts the 2,500-case corpus into a navigable structure with 12,402+ nodes, allowing stratified selection that would require months of manual review to replicate. Selection proceeds as follows: cases matching the target stratum are identified through facet pruning, ranked by analytical relevance (richness of extracted features, representativeness of the era and relevance), and reviewed against the public records request pipeline to confirm availability of supplementary redacted summaries where applicable.

3.3. Case Study Structure

Each case study is written as a structured narrative across five analytical perspectives:

1. **Platform context** — the digital environment in which the offense occurred: which platforms were used, how they were used, and what that reflects about the technological moment

2. **Perpetrator methodology** — the behavioral pattern of the offense: if contact was initiated, how grooming or sexual abuse unfolded if applicable, if there was a connection to possession of explicit material or other crimes, and what the offense type reveals about the era

3. **Investigative approach** — how law enforcement identified, investigated, and built the case: investigation type, agency network, digital forensics methods employed

4. **Prosecutorial outcomes** — charges, conviction, and sentencing, examined for what they reveal about evidentiary standards and prosecutorial priorities of the period

5. **Era relationship** — how the case reflects, complicates, or extends the patterns characteristic of its technological era, and how it connects to the other 19 case studies

3.4. Source Materials and Legal Grounding

All case studies draw exclusively from public, already-redacted records. The primary source corpus is the 2,500-case CaseLinker dataset, drawn from publicly available ICAC press releases and annual task force reports. For cases where supplementary redacted

summaries add analytical depth, records will be obtained through applicable state open records statutes:

- Arizona: Public Records Law (A.R.S. § 39-121)
- Georgia: Open Records Act (O.C.G.A. § 50-18-70 et seq.)
- Florida: Public Records Law (F.S. § 119.01)
- Texas: Public Information Act (Tex. Gov't Code § 552)
- California: California Public Records Act (Cal. Gov't Code § 6250 et seq.)
- New York: Freedom of Information Law (N.Y. Pub. Off. Law § 84 et seq.)
- Illinois: Freedom of Information Act (5 ILCS 140)

All records requests are limited to closed, adjudicated cases to minimize exemption risk under applicable state law. No attempt will be made to re-identify individuals, reconstruct victim or perpetrator identity beyond what appears in published records, or supplement public documents with external investigation.

This research received a determination from the University of Massachusetts Amherst Human Research Protection Office (HRPO Determination #7668); it does not involve private or identifiable information under federal regulations [45 CFR 46.102(f)(1), (2)].

4. CaseLinker as Research Infrastructure

4.1. From Tool to Infrastructure

The first three reports in this series documented CaseLinker as a system: its architecture, its extraction methodology, its clustering and triage logic, and its analytical outputs. For the case study work in this report and those that follow, the relevant framing is different. CaseLinker is not just a tool for this analysis—it is the research infrastructure that makes this analysis possible.

Without CaseLinker, selecting 20 representative cases from a corpus of 2,500 across four eras and five analytic dimensions would require a researcher to read 2,500 case summaries manually. The psychological cost of that—repeatedly processing disturbing case

material to identify structural patterns—is documented: digital forensics investigators in child exploitation units experience secondary traumatic stress at rates of 40–60% [12, 11]. CaseLinker's ten-dimensional facet tree converts that burden into a navigable query. A researcher identifies candidate cases in minutes rather than weeks, and without repeated full-text exposure to the most severe case material.

The broader point matters for how this work should be understood. The case studies are not the product of curated ICAC cases. They are the product of a system that has processed 2,500 cases, extracted domain related features, built a navigable corpus, and can surface precisely the cases that represent each era and dimension—cases a manual search could not reliably find and would exact a serious human cost to attempt.

4.2. Extensibility and Replication

A deliberate design principle in CaseLinker from its initial release is extensibility: the architecture is modular, the dataset is open, and the case study framework established here is replicable.

What this means concretely: the 2,500-case corpus and the extraction pipeline are released as open-source software. An investigator at an ICAC task force, a researcher at NCMEC, or an independent journalist covering digital harm can apply the same case study framework to cases in this corpus that are not covered in the 20 case studies here—extending the historical record without reproducing the infrastructure from scratch. The stratified sampling methodology, the five-dimension analytical structure, and the legal framework for records requests are all documented in this paper and applicable to additional case selection.

The goal is not a closed study with 20 conclusions. It is a replicable methodology and an open corpus that others can extend.

4.3. Vicarious Trauma and Domain Considerations

This research operates in a domain that can impose real psychological costs on everyone who engages with it seriously. Digital forensics investigators in child exploitation units experience secondary traumatic stress at rates of 40–60% [12, 11]. Researchers are not exempt. An obligation to the subject matter includes an obligation to readers and to oneself: not to unnecessarily expose material or details that causes harm without analytical purpose. The case studies in this series present structured narratives across five analytical dimensions — platform context, perpetrator methodology, investigative approach,

prosecutorial outcomes, and era relationship. They do not reproduce graphic case details, dwell on the specifics of individual offenses beyond what the analytical argument requires, or present material designed to provoke outrage rather than understanding. The goal is insight, not exposure. This extends to how perpetrators are characterized. It is tempting — and common — to reduce offenders in this domain to moral categories: monsters, predators, evil. That framing is understandable. It is also analytically unhelpful. Understanding how this crime has evolved, how perpetrators have adapted and potentially found communities within platform changes, and how law enforcement has responded requires treating cases as evidence of a systemic problem, not occasions for individual condemnation. The case studies are forward-facing: what do these cases reveal about the technological moment, the enforcement response, and the patterns that inform future prevention? That is the question. Individual offenders are not the unit of analysis. The landscape is.

5. Ethical and Dissemination Framework

5.1. Research Ethics and HRPO Determination

The ethical grounding of this research rests on three principles that have been consistent across the CaseLinker series: exclusive use of public, already-redacted records; no attempt to re-identify individuals; and thorough review prior to any public release.

HRPO Determination #7668 from the University of Massachusetts Amherst Human Research Protection Office establishes that this research does not involve private or identifiable information under federal human subjects regulations [45 CFR 46.102(f)(1), (2)]. The determination reflects the nature of the source material—closed, adjudicated cases published by law enforcement agencies for public record—not a judgment that the subject matter is without sensitivity. The sensitivity of the subject matter is precisely why the methodological choices documented here matter.

No case study will include information that risks identification of victims, witnesses, or individuals not named in original published records. Public records include names of convicted offenders and those names are available through public court records and press releases; however their names will not be included in these case studies. Cases involving the most severe severity indicators, or those whose details could compromise ongoing investigations or reveal non-public law enforcement methodology, are further abstracted

or withheld from the public release.

5.2. Restricted-First Dissemination

Case study findings will be shared with practitioner and research contacts prior to public release. This is not a legal requirement—the records are public and the research has received an HRPO not-human-subjects determination. It is a professional and domain-appropriate decision. Investigators, task force commanders, and child safety researchers who have spent careers in this domain are better positioned than any automated system to identify whether a particular case study, even derived from public records, carries risks that aggregate review might miss.

This approach also reflects the collaborative orientation of the broader CaseLinker project. The practitioners, collaborators, and mentors who have engaged with this work have been invaluable partners in validating the system's utility and appropriate scope. Restricted-first dissemination maintains that relationship.

5.3. Goal of This Work

Twenty case studies cannot capture the full history of internet-facilitated crimes against children. That is not their purpose. Their purpose is to make this issue visible, human, and actionable in ways that aggregate statistics cannot.

The officers and commanders who have spent careers in this work — building cases from digital fragments, coordinating across agencies, making triage decisions under pressure with limited resources — have accumulated institutional knowledge that has never been systematically documented at scale. The platforms and technology organizations that have deployed detection infrastructure, shared hashes across competitors, and cooperated with law enforcement have made real and measurable contributions to child safety that deserve to be understood, not just counted. Twenty structured case studies across four technological eras can begin to show what that work actually looked like in practice: the investigations that succeeded, the methodologies that evolved, the agency networks that formed around a common purpose.

But the goal of this work is not only historical documentation. It is to lower the barrier to engagement for everyone who is not already in this space. A researcher who has never read an ICAC case report, a policy analyst who has never spoken with a task force commander, a legislator who knows the issue exists but has never seen what an investigation actually involves — these are the people who need to understand this landscape to address

it effectively. Case studies are the form of evidence that reaches them. A well-constructed case study from 2013 showing how a perpetrator used early Facebook contact to initiate grooming, how law enforcement built the case, and how the prosecution succeeded tells a story that a platform involvement percentage cannot. It invites engagement from researchers, technologists, advocates, and policymakers who have something to contribute but no entry point into a domain that can feel ever-evolving, opaque, and overwhelming.

The goal, ultimately, is more people working on this problem with better information. The case studies are an invitation.

5.4. What This Research Does Not Do

Clarity about limits matters, especially in a domain where overstatement has real consequences. The case study framework established here does not:

- Re-identify or further identify victims, offenders, or witnesses beyond what appears in source publications
- Reconstruct operational law enforcement methodology beyond the source material
- Make statistical generalizations beyond what the corpus and sampling design support
- Supplement public records with external investigation, commercial data, or non-public sources
- Claim to represent the full population of ICAC investigations—only the publicly available, adjudicated subset

The adversarial risk argument formalized in Report #3 applies here: a dataset of successful prosecutions contains no failure modes, no operational tradecraft, and no evasion-relevant intelligence that is not already available through manual review of the source documents. The utility asymmetry between defenders and adversaries does not change with the addition of structured case studies.

6. Roadmap and Next Steps

6.1. Case Study Execution Timeline

This report establishes the framework. Case studies will be developed over the summer of 2026, following the methodology documented in Section 3. The execution sequence is:

1. Finalize stratified case selection using CaseLinker facet tree navigation across all four eras and five analytic dimensions
2. Submit open records requests for supplementary redacted summaries under applicable state statutes, limited to closed adjudicated cases
3. Draft case studies across the five analytical dimensions
4. Practitioner review prior to public release
5. Public release as Report #5 or as a standalone case study document, depending on volume and scope of completed studies

6.2. Parallel CaseLinker Development

The case study work runs in parallel with continued CaseLinker development. Planned additions for this period include:

- **Triage scoring refinement** — extending the priority triage system to surface cases most relevant to specific investigative contexts, not just severity
- **Temporal trend analysis** — structured tracking of platform involvement, offense type composition, and investigation type distribution across the four eras
- **Cross-era pattern detection** — identifying cases that bridge eras or exhibit methodologies characteristic of transitional periods
- **Expanded geographic coverage** — additional task force sources to increase jurisdictional diversity and reduce single-source reporting biases

These additions serve the case study work directly: richer analytical infrastructure means better-grounded case selection and more defensible pattern claims.

6.3. Community Engagement and Extension

The framework and corpus are open. Researchers, journalists, ICAC practitioners, and child safety organizations who want to apply this methodology to additional cases, extend the corpus, or develop their own analytical questions from the existing dataset are encouraged to engage directly. The open records request pipeline, the extraction pipeline, and the five-dimension case study structure are all documented and replicable. The analytical infrastructure that took months to build does not need to be rebuilt by the next person who wants to study this history.

7. Conclusion

The history of internet-facilitated crimes against children across decades of technological change exists. It is documented in over 2,500 publicly available case reports, fragmented across jurisdictions, sitting in annual reports and press releases that no one has connected at scale. The patterns and institutional record is there. The evolution is there.

What has been missing is the infrastructure to read it systematically—and the framework to turn that systematic reading into grounded historical case studies that practitioners, researchers, journalists, and the public can actually use.

This report provides that framework.

The four eras are not a theoretical construct. They reflect a real transformation in the technological landscape of exploitation and enforcement, visible in the CaseLinker corpus and documented across 16 years of public ICAC records. The 20 case studies will trace how perpetrators adapted, how platforms became vectors, how task forces responded, and what the prosecutorial record looks like when examined across eras rather than read in isolation.

The framework is also not closed. Every methodological choice documented here and in previous reports—the stratified sampling approach, the five-dimension analytical structure, the legal framework for records requests, the ethical grounding—is designed to be replicable by others. The hope is that this framework lowers the barrier for anyone who wants to engage with this domain seriously: the next researcher does not have to rebuild the infrastructure, a journalist has a platform to ask complex questions on these cases, and that the people who want to understand this landscape at scale have a pipeline that

works. The work of protecting children from online harm is too important and too large for any single system or researcher. The more people who can engage with this history rigorously, and without unnecessary exposure to its worst material — the better. That is what open infrastructure is for.

The case studies will be difficult to write. The subject matter is serious, the cases are real, and the history being reconstructed involves real harm to real children. That is precisely why the work matters—and why it requires the methodological care, moral grounding, and practitioner collaboration that this framework is designed to support.

The analysis begins this summer.

CaseLinker is available at github.com/mrinaalr/CaseLinker under the MIT License.
Live demo: web-production-13a2.up.railway.app.

References

[1] National Center for Missing & Exploited Children. (2024). *2024 CyberTipline Data Report.* https://www.missingkids.org/gethelpnow/cybertipline/cybertiplinedata

[2] Office of Juvenile Justice and Delinquency Prevention. (2024). *Internet Crimes Against Children Task Force Program.* https://ojjdp.ojp.gov/programs/internet-crimes-against-children-task-force-program

[3] Souras, Y. (2025, March 26). Testimony before the U.S. House Committee on Energy and Commerce regarding CyberTipline reporting and the REPORT Act. https://www.congress.gov/119/meeting/house/118066/witnesses/HHRG-119-IF17-Wstate-SourasY-20250326.pdf

[4] Internet Watch Foundation. (2024, July). *How AI is being abused to create child sexual abuse imagery: Updated report.* https://www.iwf.org.uk/about-us/why-we-exist/our-research/how-ai-is-being-abused-to-create-child-sexual-abuse-imagery/

[5] eSafety Commissioner. (2026, January 8). *Generative AI and child safety: A convergence of innovation and ex-*

ploitation. https://www.esafety.gov.au/newsroom/blogs/generative-ai-and-child-safety-a-convergence-of-innovation-and-exploitation

[6] Technology Coalition. (2023). *An update on voluntary detection of CSAM*. https://technologycoalition.org/resources/update-on-voluntary-detection-of-csam/

[7] Thorn. (2023, September). *Hashing and matching is core to proactive CSAM detection*. https://safer.io/resources/hashing-and-matching-is-core-to-proactive-csam-detection/

[8] Steinebach, M. (2023). An analysis of PhotoDNA. *ARES 2023 (18th International Conference on Availability, Reliability and Security)*. https://dl.acm.org/doi/fullHtml/10.1145/3600160.3605048

[9] Deryck, M., Leblanc-Albarel, D., & Preneel, B. (2026). White-box attacks on PhotoDNA perceptual hash function. *IACR ePrint*, 2026/486. https://eprint.iacr.org/2026/486.pdf

[10] United Nations Interregional Crime and Justice Research Institute. (2022). *AI for Safer Children Global Hub*. https://unicri.org/topics/AI-for-Safer-Children

[11] Burns, C. M., Morley, J., Bradshaw, R., & Domene, J. (2008). The emotional impact on and coping strategies employed by police teams investigating internet child exploitation. *Traumatology*, 14(2), 20–31.

[12] Perez, L. M., Jones, J., Englert, D. R., & Sachau, D. (2010). Secondary traumatic stress and burnout among law enforcement investigators exposed to disturbing media images. *Journal of Police and Criminal Psychology*, 25(2), 113–124.

[13] Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.

[14] Ramachandran, M. (2026). *CaseLinker: An open-source system for cross-case analysis of Internet Crimes Against Children reports. Technical Report #1*. University of Massachusetts Amherst. https://doi.org/10.48550/arXiv.2603.18020

[15] Ramachandran, M. (2026, March). *CaseLinker: Interpretable ML Approaches for Analyzing Internet Crimes Against Children Reports. Technical Report Series #2*. University of Massachusetts Amherst.

[16] Ramachandran, M. (2026, April). *CaseLinker: 5 Sources, 500 Cases, and Scaling Considerations. Technical Report Series #3*. University of Massachusetts Amherst.

[17] Ramachandran, M. (2026, March). *Why Internet Crimes Against Children And Retrospective Case Analysis Matters*. University of Massachusetts Amherst. https://mrinaalr.github.io/website/Why%20Internet%20Crimes%20Against%20Children%20And%20Retrospective%20Case%20Analysis%20Matters.pdf