

A Comparative Survey of Static and Machine Learning-Based Approaches for Detecting Malware in Multimedia Files

Shreyas Santosh Walake
Department of Computer
Engineering
JSPM's JSCOE, Pune

Prasad Vilas Wakchaure
Department of Computer
Engineering
JSPM's JSCOE, Pune

Vaishnavi Sidram Shinde
Department of Computer
Engineering
JSPM's JSCOE, Pune

Shravani Chittaranjan Takale
Department of Computer Engineering
JSPM's JSCOE, Pune

Prof. Shikha Dwivedi
Department of Computer Engineering
JSPM's JSCOE, Pune

Abstract - The exploitation of multimedia files as a delivery mechanism for malicious payloads has emerged as a significant and underaddressed threat vector, driven by the widespread exchange of images, video, audio, and document files through messaging platforms and email. Unlike executable-based malware, threats embedded within media containers exploit steganographic techniques and structural format manipulation to evade conventional signature-based detection, which remains predominantly optimized for executable file inspection. This paper presents a comparative survey of existing detection approaches for media-borne malware, spanning signature-based antivirus engines, rule-based heuristic and static analysis frameworks, classical machine learning techniques including chi-square and Regular-Singular steganalysis, and deep learning architectures based on convolutional neural networks. Each category of approach is examined with respect to detection accuracy, computational resource requirements, format coverage, deployment constraints, and resilience against adaptive steganographic algorithms such as J-UNIWARD. A structured comparative analysis reveals that no existing approach simultaneously achieves multi-format coverage, offline deployability and robustness against adaptive embedding techniques, with deep learning methods offering superior detection of adaptive steganography at the cost of GPU dependency and interpretability, while classical and heuristic methods remain CPU-deployable but format-fragmented and limited against adaptive algorithms. Based on this analysis, the paper identifies four principal research gaps in the current literature and outlines a research direction characterized by hybrid heuristic-statistical and machine learning fusion as a viable path toward unified, resource-efficient, multi-format detection. This survey is intended to consolidate fragmented research across the steganalysis and malware detection communities and to inform future system design choices in this domain.

Indexed Terms - Malware Detection, Multimedia Security, Steganography, Steganalysis, Static Analysis, YARA Rules, Random Forest, Hybrid Detection Framework, Digital Forensics, Offline Security

I. INTRODUCTION

Multimedia files have become one of the most widely exchanged forms of digital content, transmitted in vast volumes daily through messaging platforms, email, and cloud storage services. This scale of exchange has attracted the attention of threat actors, who increasingly exploit image, video, audio, and document files as carriers for malicious payloads. Unlike conventional executable-based malware, threats embedded within multimedia containers exploit format-specific structural properties and steganographic embedding techniques to conceal malicious content without altering the perceptible characteristics of the carrier file, allowing them to bypass detection mechanisms designed primarily for executable inspection.

The research response to this threat has evolved across multiple, largely disconnected directions. Signature-based antivirus engines remain the dominant commercial defense but are architecturally limited to known-pattern matching and demonstrate poor coverage of format-internal anomalies [15]. Classical steganalysis research has produced statistical detection methods such as chi-square analysis [1] and Regular-Singular estimation [2], both effective against simple embedding but documented to fail against natural image complexity [3] and adaptive algorithms [4]. More recent work has shifted toward convolutional neural network architectures [5], [11], [12], [18], [19] capable of detecting adaptive steganography such as J-UNIWARD [4] at the cost of substantial computational resources and reduced interpretability. Parallel to this, behavioral and dynamic analysis research addresses runtime threats that static methods

cannot observe [16], while heuristic and rule-based static analysis frameworks offer lightweight, interpretable detection at the cost of vulnerability to novel obfuscation patterns.

Each of these research directions has been evaluated largely in isolation, within its own benchmark conditions and threat assumptions, with limited cross-comparison of accuracy, computational cost, format coverage, and deployment feasibility under a common framework. Practitioners and researchers seeking to design a detection system for media-borne malware are consequently required to synthesize findings across steganalysis literature, malware detection surveys, and commercial security documentation that rarely engage with one another directly.

This paper addresses that gap through a structured comparative survey of existing detection approaches for malware concealed within multimedia files. The survey is organized around four major technique families: signature-based and commercial antivirus detection, rule-based heuristic and static analysis frameworks, classical machine learning-based steganalysis, and deep learning-based steganalysis. Each family is examined with respect to five comparison dimensions: detection accuracy, computational resource requirements, file format coverage, deployment constraints including offline capability, and resilience against adaptive steganographic embedding.

The contributions of this survey are threefold. First, a consolidated comparative analysis of detection approaches across signature-based, heuristic, classical machine learning, and deep learning categories, evaluated under a common set of comparison dimensions rather than within isolated benchmark contexts. Second, identification of four principal research gaps that persist across the surveyed literature, including the absence of any approach achieving simultaneous multi-format coverage, offline deployability, and adaptive steganography resilience. Third, a proposed research direction characterized by hybrid fusion of heuristic and machine learning detection, motivated directly by the comparative weaknesses identified across individual technique families.

II. LITERATURE SURVEY

A. Signature-Based and Commercial Antivirus Detection

Signature-based detection remains the most widely deployed defense mechanism in commercial antivirus products, relying on known byte-pattern matching against a maintained database of malicious file hashes and code signatures. This approach is computationally efficient and produces near-zero false positives against previously catalogued threats, but its detection capability is fundamentally bounded by the contents of its signature database. Idika and Mathur [15] characterize this limitation as inherent to the technique rather than an implementation deficiency: signature matching cannot identify threats that

have not been previously observed and catalogued, regardless of how the underlying malicious behavior is expressed. This limitation is particularly acute for media-borne threats, where the malicious payload is often a generic dropper or beacon stub repackaged behind a novel steganographic or container-injection technique specific to the carrier file format. Commercial antivirus engines, optimized primarily for executable file scanning, frequently lack the format-specific parsing logic required to inspect the internal structure of image, video, audio, and document containers, leaving a substantial inspection gap that signature matching alone cannot close.

B. Rule-Based Heuristic and Static Analysis Frameworks

Heuristic and static analysis approaches extend beyond signature matching by evaluating structural and statistical properties of a file without executing it. Rule-based frameworks such as YARA enable pattern-based detection across arbitrary byte sequences and have been widely adopted for identifying embedded executable signatures, suspicious script content, and known malware family indicators within non-executable carrier files. Complementary statistical techniques, including Shannon entropy analysis and container overlay detection, identify anomalous structural properties such as appended data beyond declared format boundaries or abnormal randomness inconsistent with a file's expected compression characteristics. These approaches offer the advantage of interpretability, low computational overhead, and CPU-only deployability, making them suitable for resource-constrained and offline environments. Their principal weakness is coverage fragility: heuristic rules and statistical thresholds are inherently reactive to known attack patterns and structural anomalies, and a sufficiently novel obfuscation or encoding scheme can evade detection without triggering any existing rule, requiring continuous rule-base maintenance as new evasion techniques emerge.

C. Classical Machine Learning-Based Steganalysis

Classical steganalysis research predates deep learning approaches and remains relevant due to its computational efficiency and interpretability. Westfeld and Pfitzmann [1] established chi-square analysis as a statistical test for detecting least-significant-bit embedding, based on the observation that LSB substitution produces measurable uniformity in bit-pair frequencies inconsistent with natural image or audio statistics. Fridrich et al. [2] introduced Regular-Singular analysis, providing a quantitative estimator of LSB embedding rate that remains a standard benchmark comparison point in subsequent steganalysis literature. Both techniques perform reliably against naive LSB embedding but are documented by Ker [3] to fail against natural images and audio recordings exhibiting high intrinsic statistical complexity, where sensor or recording noise produces LSB distributions statistically indistinguishable from genuine embedding. This limitation

becomes more pronounced against adaptive embedding algorithms such as J-UNIWARD [4], which deliberately minimize the statistical footprint of embedding by targeting complex image regions, rendering classical statistical tests substantially less effective. Random Forest and other ensemble methods applied to hand-crafted statistical features [8] offer improved detection of tool-based steganography with fixed statistical signatures but inherit similar limitations against adaptive embedding, since hand-crafted features designed around known embedding signatures do not generalize to algorithms engineered specifically to avoid producing such signatures.

D. Deep Learning-Based Steganalysis

The limitations of hand-crafted statistical features against adaptive steganography motivated a shift toward convolutional neural network architectures capable of learning discriminative representations directly from image data. Qian et al. [19] demonstrated that CNNs can learn steganalytic features without manual engineering, establishing the feasibility of the deep learning approach. Xu et al. [11] subsequently refined CNN architectural design specifically for steganalysis, while Ye et al. [12] introduced hierarchical feature representations that improved detection accuracy against adaptive algorithms. Yedroudj et al. [18] advanced this direction further with a computationally efficient spatial-domain CNN achieving competitive accuracy at reduced training cost. Boroumand et al. [5] established a strong benchmark with SRNet, a deep residual architecture achieving 64.3% accuracy on ALASKA2 J-UNIWARD detection, while Yousfi et al. [9] demonstrated that transfer learning from pre-trained CNN feature extractors via EfficientNet-B4 achieves 68.1% accuracy on the same benchmark, representing among the strongest published results for this specific adaptive embedding algorithm. These results collectively establish that meaningful detection of adaptive steganography requires CNN architectures with substantial parameter counts and correspondingly GPU-dependent training and inference, a resource requirement that constrains deployability in offline, CPU-only, or air-gapped environments where lightweight detection is operationally necessary.

E. Behavioral and Dynamic Analysis Approaches

A separate research direction addresses threats that static analysis, regardless of sophistication, cannot observe: malicious behavior that only manifests at runtime. Sikorski and Honig [16] document dynamic analysis techniques including sandboxed execution monitoring, which observe actual program behavior such as network connections, process injection, and file system modification rather than inferring intent from static structural properties. This approach is particularly relevant to media-borne threats that rely on macro execution, embedded script interpreters, or document-viewer exploit chains, where the malicious

action only occurs once the carrier file is opened in a vulnerable application context. Dynamic analysis offers detection coverage that static methods structurally cannot provide, but requires an isolated execution environment, introduces non-trivial latency per file, and is unsuitable as a pre-screening mechanism for high-volume file exchange scenarios such as messaging platform attachments, where files must be evaluated before being opened rather than after.

F. Summary of Surveyed Categories

Across the five categories surveyed, a consistent pattern emerges: techniques that offer strong detection capability against sophisticated or adaptive threats tend to require substantial computational resources or execution-time observation, while techniques that are lightweight and deployable in constrained environments tend to be limited to known patterns or non-adaptive embedding signatures. No single surveyed category resolves this tradeoff independently, motivating the comparative analysis presented in Section III.

III. COMPARATIVE ANALYSIS

This section presents a structured comparison of the detection approaches surveyed in Section II across five dimensions identified as critical for practical deployment against media-borne malware: detection accuracy against the threat type the approach targets, computational resource requirements, file format coverage, offline deployment feasibility, and resilience against adaptive steganographic embedding. Table I summarizes this comparison across the five technique categories.

TABLE I: COMPARATIVE ANALYSIS OF MEDIA-BORNE MALWARE DETECTION APPROACHES

Approach	Detection Accuracy	Compute Requirement	Format Coverage	Offline Capable	Adaptive Stego Resilience
Signature-based AV [15]	High on known threats; near-zero on novel payloads	Low (CPU)	Low — executable-focused	Yes	None
Heuristic/Static (YARA, entropy)	High on structural anomalies	Low (CPU)	High — format-agnostic rules	Yes	Low

Approach	Detection Accuracy	Compute Requirement	Format Coverage	Offline Capable	Adaptive Stego Resilience
	es; format-dependent				
Classical ML (chi-square [1], RS [2], RF [8])	High on non-adaptive embedding; near-random on adaptive	Low (CPU)	Format-specific per model	Yes	Low
Deep Learning (CNN/SRNet [5], EfficientNet [9])	Highest on adaptive embedding (64–68%)	High (GPU)	Format-specific per model	Limited	High
Behavioral/Dynamic [16]	High on runtime threats; blind to dormant payloads	Moderate–High	Format-agnostic	Limited	N/A (different threat class)

Table I. Comparative summary across five detection technique families. Adaptive steganography resilience reflects documented performance specifically against algorithms such as J-UNIWARD [4] that minimize embedding distortion footprints.

A. Accuracy Versus Resource Tradeoff

The comparison surfaces a consistent inverse relationship between detection accuracy against sophisticated threats and resource efficiency. Signature-based and heuristic approaches achieve high accuracy only against threats matching pre-defined patterns, with accuracy degrading sharply against novel or adaptively obfuscated payloads. Deep learning approaches invert this relationship, achieving the strongest documented results against adaptive steganography [5], [9] but requiring GPU inference infrastructure that is frequently unavailable in the offline, air-gapped, or resource-constrained deployment contexts where media-borne threats are most prevalent.

B. Format Coverage Fragmentation

A second pattern emerges in format coverage. Heuristic and rule-based frameworks are the only category achieving broad format-agnostic coverage, since structural checks

such as entropy analysis and overlay detection generalize across container types with format-specific parameter tuning. Classical and deep learning steganalysis models, by contrast, are typically trained and evaluated on a single format — JPEG, PNG, or similar — with no surveyed work demonstrating a single model or framework achieving comparable accuracy across image, video, audio, and document formats simultaneously. This fragmentation means that comprehensive multi-format protection currently requires assembling multiple independently trained models rather than relying on any single surveyed approach.

C. Deployment Feasibility

Offline deployability favors signature-based and heuristic approaches, both of which operate entirely on local computation without external dependency. Deep learning approaches are constrained by GPU requirements that are difficult to satisfy in air-gapped or low-resource environments, while behavioral and dynamic analysis approaches require an isolated sandboxed execution environment that introduces both infrastructure complexity and per-file latency unsuitable for high-throughput pre-screening scenarios such as messaging platform attachment scanning.

D. Adaptive Steganography as the Unresolved Boundary

Resilience against adaptive steganographic embedding represents the most consistent boundary across the surveyed literature. With the exception of deep learning approaches, every surveyed technique category demonstrates substantial accuracy degradation against algorithms such as J-UNIWARD [4] that are specifically engineered to minimize detectable statistical footprint. Even within deep learning approaches, the strongest documented results [5], [9] remain in the 64–68% accuracy range — a meaningful improvement over near-random classical performance, but far from the near-perfect accuracy achievable against non-adaptive embedding methods. This indicates that adaptive steganography resistance remains a partially, rather than fully, solved problem even within the most resource-intensive surveyed approach.

IV. RESEARCH GAPS AND LIMITATIONS

The comparative analysis in Section III surfaces four research gaps that persist across the surveyed literature, each arising directly from a tradeoff pattern identified in the preceding comparison.

A. Absence of Simultaneous Multi-Format, Offline, and Adaptive-Resilient Detection

No surveyed approach achieves all three properties simultaneously. Heuristic and signature-based methods achieve format breadth and offline deployability but fail against adaptive embedding. Deep learning methods achieve adaptive resilience but sacrifice offline deployability due to

GPU dependency. This three-way tradeoff, evident across Table I, indicates that existing research has optimized within individual technique families rather than across the combined constraint space that real-world deployment — particularly in messaging platform pre-screening contexts — actually requires.

B. Fragmented Evaluation Methodology Across the Literature

Surveyed steganalysis and detection approaches are predominantly evaluated within single-format, single-dataset benchmarks specific to each technique family, with limited cross-study comparison under shared evaluation conditions. This fragmentation, evident in the difficulty of directly comparing accuracy figures across Sections II-C and II-D without first normalizing for dataset, format, and embedding algorithm differences, complicates objective assessment of which approach is genuinely preferable for a given deployment scenario rather than simply better-suited to its own benchmark.

C. Limited Treatment of Resource-Constrained and Mobile Deployment Contexts

The majority of surveyed deep learning steganalysis research targets accuracy maximization on academic benchmarks without explicit consideration of inference latency, memory footprint, or CPU-only deployment constraints. Given that media-borne malware is predominantly delivered through mobile-first messaging platforms, this represents a meaningful disconnect between where the threat is most prevalent and where the strongest detection research has been validated.

D. Underexplored Fusion of Heuristic and Learned Detection Signals

While heuristic and machine learning-based approaches are extensively studied in isolation, the surveyed literature contains comparatively limited formal treatment of hybrid architectures that combine rule-based structural detection with statistical or learned classification within a single decision framework. Existing work tends to treat these as competing rather than complementary techniques, despite their documented weaknesses being largely non-overlapping — heuristic methods struggle with novel obfuscation, while learned methods struggle with format generalization and resource cost, suggesting unrealized complementarity rather than genuine redundancy between the two approaches.

V. PROPOSED RESEARCH DIRECTION

The research gaps identified in Section IV converge on a single underlying design tension: detection approaches that generalize across file formats and remain deployable offline are consistently vulnerable to adaptive embedding

techniques, while approaches resilient to adaptive embedding are consistently constrained to single formats and dependent on GPU infrastructure. Addressing this tension does not require resolving it within a single monolithic technique, but rather through architectural composition that assigns each technique family to the threat class it is best suited to address.

A. Hybrid Heuristic-Statistical and Machine Learning Fusion

A viable research direction, motivated directly by Gap D in Section IV, involves combining rule-based heuristic detection with statistically-grounded machine learning classification within a single decision pipeline rather than treating the two as competing standalone systems. Heuristic detection, including signature matching, entropy analysis, and structural anomaly detection, can provide broad format coverage and immediate detection of known attack patterns at negligible computational cost. Machine learning classification, applied selectively using hand-crafted statistical features rather than raw learned representations, can extend detection coverage to tool-based steganographic embedding with fixed statistical signatures, such as quantization table manipulation or alpha channel encoding, without requiring GPU infrastructure. This composition directly addresses Gap A by achieving multi-format coverage and offline deployability simultaneously, though it does not independently resolve resilience against adaptive embedding algorithms such as J-UNIWARD, which remains a boundary condition for any architecture relying on hand-crafted rather than learned representations.

B. Tiered Detection Architecture for Resource-Constrained Deployment

Addressing Gap C requires explicit architectural accommodation for resource-constrained and mobile-adjacent deployment contexts, rather than treating computational efficiency as a secondary optimization applied after accuracy maximization. A tiered architecture, in which lightweight heuristic and classical statistical detection operate as a first-pass filter on resource-constrained devices, with computationally intensive deep learning analysis reserved for an optional second-pass tier invoked only on files flagged as uncertain by the first tier, offers a path toward reconciling the accuracy-resource tradeoff identified in Section III-A without requiring every file to incur the full computational cost of adaptive-steganography-resilient detection.

C. Standardized Cross-Technique Evaluation Protocols

Addressing Gap B requires the establishment of shared evaluation benchmarks spanning multiple file formats and

embedding algorithms, against which heuristic, classical machine learning, and deep learning approaches can be assessed under identical conditions. The absence of such standardized protocols in the current literature, evident in the

format- and dataset-specific nature of the benchmarks surveyed in Section II, limits the field's ability to make objective comparative claims about which technique family is preferable for a given deployment scenario, independent of vendor or research-group benchmark selection.

D. Outlook

The convergence of these three directions — hybrid detection fusion, tiered resource-aware architecture, and standardized evaluation — does not eliminate the fundamental tradeoff between adaptive steganography resilience and offline deployability identified throughout this survey, but offers a structured path toward narrowing it. Future research validating such architectures against the specific evaluation protocols proposed in Section V-C would provide the field with comparative evidence currently absent from the fragmented benchmark landscape described in Section IV-B.

VI. CONCLUSION

This paper presented a comparative survey of detection approaches for malware concealed within multimedia files, spanning signature-based antivirus engines, rule-based heuristic and static analysis frameworks, classical machine learning steganalysis, deep learning-based steganalysis, and behavioral dynamic analysis. The comparative analysis across five dimensions — detection accuracy, computational requirement, format coverage, offline deployability, and adaptive steganography resilience — revealed a consistent tradeoff pattern: techniques achieving broad format coverage and offline deployability are vulnerable to adaptive embedding algorithms such as J-UNIWARD [4], while techniques resilient to adaptive embedding, achieving 64–68% accuracy in the strongest documented cases [5], [9], require GPU infrastructure that constrains deployment in offline and resource-constrained contexts.

Four research gaps were identified from this analysis: the absence of any single approach achieving multi-format coverage, offline deployability, and adaptive resilience simultaneously; fragmented evaluation methodology across format- and dataset-specific benchmarks; limited treatment of mobile and resource-constrained deployment contexts despite their relevance to the threat's primary delivery channels; and underexplored fusion between heuristic and machine learning detection signals despite their largely non-overlapping weaknesses. In response, this paper outlined a research direction centered on hybrid heuristic-statistical and machine learning fusion, tiered resource-aware architecture, and standardized cross-technique evaluation protocols, while explicitly acknowledging that this direction

narrows rather than eliminates the adaptive steganography resilience boundary identified throughout the survey. By consolidating findings that are currently fragmented across steganalysis, malware detection, and dynamic analysis research communities, this survey is intended to inform future system design choices and evaluation practices in the domain of media-borne malware detection, an area of growing relevance given the continued reliance on multimedia file exchange across messaging and communication platforms.

VII. REFERENCES

- [1] A. Westfeld and A. Pfitzmann, "Attacks on steganographic systems," in *Proc. 3rd Int. Workshop Information Hiding*, vol. 1768, Springer, 1999, pp. 61–76. [Online]. Available: <https://ieeexplore.ieee.org/document/648030>
- [2] J. Fridrich, M. Goljan, and R. Du, "Reliable detection of LSB steganography in color and grayscale images," in *Proc. ACM Workshop Multimedia and Security*, Ottawa, Canada, 2001, pp. 27–30. [Online]. Available: <https://dl.acm.org/doi/10.1145/1232454.1232466>
- [3] A. D. Ker, "A general framework for the statistical steganalysis of LSB embedding," in *Proc. 9th Int. Workshop Information Hiding*, vol. 4567, Springer, 2007, pp. 296–311. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-540-77370-2_20
- [4] V. Holub, J. Fridrich, and T. Denemark, "Universal distortion function for steganography in an arbitrary domain," *EURASIP Journal on Information Security*, vol. 2014, no. 1, pp. 1–13, 2014. [Online]. Available: <https://ieeexplore.ieee.org/document/6595569>
- [5] M. Boroumand, M. Chen, and J. Fridrich, "Deep residual network for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 5, pp. 1181–1193, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8470101>
- [6] R. Cogranne, Q. Giboulot, and P. Bas, "The ALASKA steganalysis challenge: A first step towards steganalysis into the wild," in *Proc. ACM Workshop on Information Hiding and Multimedia Security*, Paris, France, 2019, pp. 1–10. [Online]. Available: <https://dl.acm.org/doi/10.1145/3335203.3335238>
- [7] P. Bas, T. Filler, and T. Pevný, "'Break our steganographic system' — the ins and outs of organizing BOSS," in *Proc. 13th Int. Workshop Information Hiding*, vol. 6958, Springer, 2011, pp. 59–70. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-642-24178-9_5
- [8] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001. [Online]. Available: <https://link.springer.com/article/10.1023/A:1010933404324>
- [9] Y. Yousfi, J. Butora, J. Fridrich, and P. Bas, "Breaking ALASKA: Color images steganalysis using a pre-trained CNN feature extractor," in *Proc. IS&T Electronic Imaging, Digital Forensics and Watermarking*, 2020, pp. 413-1–413-7. [Online]. Available: <https://doi.org/10.2352/ISSN.2470-1173.2020.4.DFWM-413>
- [10] A. Selvaraj, A. Ezhilarasan, S. L. J. Wellington, and A. Roy Sam, "Digital image steganalysis: A survey on paradigm shift from machine learning to deep learning based techniques," *IET Image Processing*, vol. 14, no. 12, pp. 2393–2411, 2020. [Online]. Available: <https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/iet-ipr.2020.0543>
- [11] G. Xu, H.-Z. Wu, and Y.-Q. Shi, "Structural design of convolutional neural networks for steganalysis," *IEEE Signal Processing Letters*, vol. 23, no. 5, pp. 708–712, May 2016. [Online]. Available: <https://ieeexplore.ieee.org/document/7444146>
- [12] J. Ye, J. Ni, and Y. Yi, "Deep learning hierarchical representations for image steganalysis," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 11, pp. 2545–2557, 2017. [Online]. Available: <https://ieeexplore.ieee.org/document/7937836>
- [13] R. Chaganti, V. Ravi, and T. D. Pham, "A deep learning based approach for detecting novel steganography techniques in digital images," *arXiv preprint arXiv:2104.05480*, 2021. [Online]. Available: <https://arxiv.org/abs/2104.05480>

- [14] F. Strachanski, N. Strodthoff, and S. Schneider, "A comprehensive pattern-based overview of stegomalware," in *Proc. ARES CUING Workshop*, Vienna, Austria, 2024. [Online]. Available: <https://dl.acm.org/doi/10.1145/3664476.3670439>
- [15] I. Idika and A. P. Mathur, "A survey of malware detection techniques," Department of Computer Science, Purdue University, Tech. Rep., 2007. [Online]. Available: <https://www.cs.purdue.edu/homes/apm/papers/>
- [16] M. Sikorski and A. Honig, *Practical Malware Analysis: The Hands-On Guide to Dissecting Malicious Software*. San Francisco, CA, USA: No Starch Press, 2012.
- [17] A. Souri and R. Hosseini, "A state-of-the-art survey of malware detection approaches using data mining techniques," *Human-Centric Computing and Information Sciences*, vol. 8, no. 1, pp. 1–22, 2018. [Online]. Available: <https://doi.org/10.1186/s13673-018-0125-x>
- [18] M. Yedroudj, F. Comby, and M. Chaumont, "Yedroudj-Net: An efficient CNN for spatial steganalysis," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, 2018, pp. 2092–2096. [Online]. Available: <https://ieeexplore.ieee.org/document/8461438>
- [19] Y. Qian, J. Dong, W. Wang, and T. Tan, "Deep learning for steganalysis via convolutional neural networks," in *Proc. SPIE Media Watermarking, Security, and Forensics*, vol. 9409, San Francisco, CA, USA, 2015. [Online]. Available: <https://doi.org/10.1117/12.2083479>
- [20] S. Lyu and H. Farid, "Steganalysis using higher-order image statistics," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 1, pp. 111–119, Mar. 2006. [Online]. Available: <https://ieeexplore.ieee.org/document/1597126>