Metadata of the chapter that will be visualized in SpringerLink

Book Title	Proceedings of 11th International Congress on Information and Communication Technology	
Series Title		
Chapter Title	Efficient Detection of Child Sexual Abuse Material in Low-Resource Forensic Environments Using Transfer Learning with Lightweight Convolutional Neural Networks	
Copyright Year	2026	
Copyright HolderName	The Author(s), under exclusive license to Springer Nature Switzerland AG	
Corresponding Author	Family Name	**Dalpian**
	Particle	
	Given Name	**Guilherme Martini**
	Prefix	
	Suffix	
	Role	
	Division	Department of Electrical Engineering
	Organization	University of Brasília
	Address	Brasília, DF, 70910-900, Brazil
	Email	guilherme.dalpian@aluno.unb.br
Author	Family Name	**Nze**
	Particle	
	Given Name	**Georges Daniel Amvame**
	Prefix	
	Suffix	
	Role	
	Division	Department of Electrical Engineering
	Organization	University of Brasília
	Address	Brasília, DF, 70910-900, Brazil
	Email	g.amvame@unb.br
Author	Family Name	**Mendonça**
	Particle	**de**
	Given Name	**Fábio Lúcio Lopes**
	Prefix	
	Suffix	
	Role	
	Division	Department of Electrical Engineering
	Organization	University of Brasília
	Address	Brasília, DF, 70910-900, Brazil
	Email	fabio.mendonca@unb.br

Abstract	The exponential growth in digital evidence has created increased backlogs in forensic investigations, including cases involving Child Sexual Abuse Material (CSAM), where timely action is paramount. To address this, we present a validated, cost-effective, and field-deployable triage solution to accelerate the identification of previously unknown illicit content on-site. The system integrates a secure, write-blocked Linux environment with a hybrid solution based on the IPED forensic framework. It combines efficient

hash database filtering (NIST CAID, ProjectVIC) for known files with a lightweight EfficientNetV2-B0 Convolutional Neural Network (CNN), optimized for CPU processing via ONNX conversion, to classify unknown material, including video frames. Evaluation on a test set of 41,931 images across three classes (CSAM, porn, other) showed a global accuracy of 96.5%, with a recall of 95.8% and a precision of 91.4% for the CSAM class. Benchmarking confirms the system's efficiency on legacy hardware, with processing speeds of 30–60 images per second on standard CPUs. This open-source methodology enables law enforcement to identify unknown illicit content in near real-time, reducing analysts' exposure to traumatic material and supporting both operational efficiency and occupational health.

Keywords (separated by '-') CSAM detection - deep learning - transfer learning - CNN - digital forensics - triage

Author Proof

Check for updates

Efficient Detection of Child Sexual Abuse Material in Low-Resource Forensic Environments Using Transfer Learning with Lightweight Convolutional Neural Networks

Guilherme Martini Dalpian(✉), Georges Daniel Amvame Nze, and Fábio Lúcio Lopes de Mendonça

Department of Electrical Engineering, University of Brasília, Brasília, DF 70910-900, Brazil
guilherme.dalpian@aluno.unb.br, {g.amvame,fabio.mendonca}@unb.br

Abstract. The exponential growth in digital evidence has created increased backlogs in forensic investigations, including cases involving Child Sexual Abuse Material (CSAM), where timely action is paramount. To address this, we present a validated, cost-effective, and field-deployable triage solution to accelerate the identification of previously unknown illicit content on-site. The system integrates a secure, write-blocked Linux environment with a hybrid solution based on the IPED forensic framework. It combines efficient hash database filtering (NIST CAID, ProjectVIC) for known files with a lightweight EfficientNetV2-B0 Convolutional Neural Network (CNN), optimized for CPU processing via ONNX conversion, to classify unknown material, including video frames. Evaluation on a test set of 41,931 images across three classes (CSAM, porn, other) showed a global accuracy of 96.5%, with a recall of 95.8% and a precision of 91.4% for the CSAM class. Benchmarking confirms the system's efficiency on legacy hardware, with processing speeds of 30–60 images per second on standard CPUs. This open-source methodology enables law enforcement to identify unknown illicit content in near real-time, reducing analysts' exposure to traumatic material and supporting both operational efficiency and occupational health.

AQ1 AQ2

Keywords: CSAM detection · deep learning · transfer learning · CNN · digital forensics · triage

1 Introduction and Justification

Digital forensic laboratories face substantial challenges from rapidly increasing data volumes, leading to case backlogs and delayed investigations [3,22]. In cases involving Child Sexual Abuse Material (CSAM), immediate on-site intelligence

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
X.-S. Yang et al. (Eds.): ICICT 2026, LNNS 2120, pp. 1–10, 2026.
https://doi.org/10.1007/978-3-032-34434-2_23

Author Proof

about the presence of victims and offenders is essential [30]. Current detection methods primarily rely on hash analysis with established databases such as ProjectVIC and NIST CAID [18,19]. Although effective for identifying previously cataloged material, these methods are unable to detect new, unknown, or modified content. Relying only on hash-based tools in the field risks overlooking evidence of ongoing criminal activity or unidentified victims. Immediate detection of unfamiliar material is critical for prompt investigative response.

Another challenge is the psychological distress from repeated manual exposure to graphic content [5]. Automated tools serve a dual function by advancing investigations and providing occupational safeguards for staff.

Shavers [25] defines forensic triage as processing computer evidence at the crime scene to rapidly identify the most critical evidence before comprehensive laboratory examination, while preserving forensic integrity. Although forensic distributions such as CAINE, Tsurugi, and PALADIN are available, and AI use in forensics is expanding [6,7], a gap remains: there is no system specifically designed for on-site, open-source, low-resource detection of unknown CSAM. This study addresses this gap by introducing a lightweight, field-ready AI plugin for the open-source IPED framework that leverages a dual AI model combining hash-based pre-filtering with deep learning classification.

2 Objective

The primary objective is to design and implement an integrated framework for rapid on-site forensic analysis during crime scenes and search-and-seizure operations. The proposed tool employs a hybrid detection method, utilizing cryptographic hashes for recognized datasets and Convolutional Neural Networks to identify novel illicit material. The system targets at least 90% accuracy in detecting illegal content while maintaining inference times suitable for practical deployment, specifically less than 50 milliseconds per image on CPU-only systems.

3 Related Work and Scientific Gap

Digital forensic triage addresses the challenge of increasing file volumes by enabling rapid on-site evidence verification [14,20,21]. The use of bootable Linux distributions and custom scripts has enhanced field triage efficiency, with write-blocking and evidence segregation remaining essential [2,26].

Evidence integrity is essential for admissibility [24]. Laboratories typically rely on hardware write blockers, whereas field solutions frequently employ robust software approaches enabled by open-source tools such as Linux. Disk encryption technologies such as BitLocker pose challenges, though open-source tools like dislocker provide support when credentials are available [4].

The IPED framework is an extensible forensic analysis tool developed by the Brazilian Federal Police that includes data carving, hash set matching, and many other features [12]. Many machine learning technologies for CSAM detection

Author Proof

have shown great potential, but careful dataset management is needed to reduce false positives [8,18]. Transfer learning [15,29] with lightweight models such as MobileNet and EfficientNet enables accurate training with limited data [16,23,28].

Despite these advancements, no device currently provides fully integrated hash-based filtering and automatic AI classification for CSAM detection in a portable, open-source, field-ready, and forensically robust format. This study addresses this gap by introducing a validated hybrid approach that combines both methods within a secure, deployable Linux environment.

4 Theoretical Foundation

Convolutional Neural Networks (CNNs), a core component of modern Artificial Intelligence, work with grid-structured data, like images, by replacing traditional matrix multiplication with convolution in at least one layer. Using learnable kernels, these networks leverage parameter sharing and sparse connectivity to efficiently capture spatial feature hierarchies, reducing the preprocessing steps required by legacy algorithms. Typically, the architecture follows convolution and pooling layers for feature extraction, followed by fully connected layers for classification [17].

Transfer learning leverages knowledge from large source datasets to enhance performance on smaller target datasets. The key insight is that learned feature representations capture domain-general properties transferable across tasks [29]. Feature extraction involves freezing pre-trained layers and training only new classification layers, whereas fine-tuning unfreezes higher layers and applies reduced learning rates to specialize them. Data augmentation increases diversity through random transformations, thereby preventing overfitting and improving generalization [27]. Techniques such as dropout, batch normalization, and weight regularization further reduce model complexity [10].

Data collection in forensic scenarios with imbalanced classes, such as CSAM, presents significant challenges. Statistical metrics may be misleading, and models often rely on predictions from the majority class [9]. Mitigation strategies include weighted loss functions, stratified splitting, and balanced validation sets [11].

5 Methodology

The research uses Design Science Research methodology, combining rigorous experimentation with practical artifact development. The dataset includes 395,340 images in three forensically relevant classes: CSAM (46,662 images from Brazilian Federal Police known-illegal archives), Porn (130,933 NSFW images from public datasets), and Other (217,745 benign images, system icons, and software files). All CSAM handling followed strict ethical protocols: authorized personnel only, air-gapped processing, encrypted storage, separation of derivative work, and full compliance with the LGPD [1].

Author Proof

Data preprocessing included SHA-1 hashing to eliminate exact duplicates and verifying image file integrity. To prevent train-test contamination, where visually similar images appear in both sets, embeddings were generated from a pre-trained CNN to identify visually similar groups, and entire clusters were assigned to distinct sets (train, validation, or test). The validation dataset was balanced to ensure the majority class did not overwhelm minority classes. The final dataset distribution is presented in Table 1.

Table 1. Dataset Distribution

Class	Train	Validation	Test	Total
csam	32,734	9,295	4,633	46,662
porn	108,555	9,295	13,083	130,933
other	184,235	9,295	24,215	217,745
Total	325,524	27,885	41,931	395,340

EfficientNetV2B0 [28] was selected for its optimal balance between accuracy and efficiency, offering a favorable tradeoff among ImageNet accuracy (78.7%), model size (23 MB), and inference speed. These characteristics make it well-suited for deployment on resource-constrained devices.

After dataset preparation and model selection, we implemented a two-stage training approach. Stage 1 (Feature Extraction, 50 epochs) froze all convolutional layers of EfficientNetV2B0 and added a new classification head with global average pooling, dropout (0.4), and an output layer (3 units, softmax), training only these new classification layers with a learning rate of 1×10^{-3}. Stage 2 (Fine-tuning, 50 epochs) unfroze the final 20% of the layers and applied a conservative learning rate 1×10^{-5} to prevent catastrophic forgetting while enabling task-specific specialization.

Given the dataset's imbalance, we calculated class weights proportional to the number of training samples per class. The weighting scheme employed the formula $w_i = \frac{\text{total_samples}}{N_{\text{classes}} \times \text{count}_i}$, where total_samples $= 325,524$, $N_{\text{classes}} = 3$, and count_i represents the sample count for class i. This approach inherently penalizes errors on minority classes (CSAM and porn) while reducing the penalty on the majority class (other). Furthermore, an optional multiplicative parameter was explicitly applied to the CSAM class weight to balance between forensic precision (false positives) and detection recall (false negatives). In the training configuration chosen for this work, a multiplicative factor of 0.8 was applied to the base CSAM weight, resulting in final class weights of csam: 2.652, porn: 1.000, and other: 0.589. This conservative weighting prioritizes forensic precision—minimizing false accusations while maintaining high recall for unknown CSAM.

To ensure compatibility in CPU-only field environments, the trained model was converted to ONNX (Open Neural Network Exchange) format using the tf2onnx library in 32-bit floating-point, enabling inference without heavy TensorFlow or PyTorch dependencies while maintaining accuracy [13]. ONNX Run-

Author Proof

time provides optimized CPU inference, and the final model size of 23 MB makes it portable via USB.

5.1 Forensic Linux Environment Implementation

The proposed system integrates a secure, write-blocked Linux environment, enabling both on-site and laboratory deployment. The bootable ISO environment, derived from Kali Linux, can be executed directly on the machine to be examined by booting from a USB drive during field operations, or by an examiner accessing a removed hard drive from another machine. The Live-Linux environment operates entirely in RAM, ensuring no writes to examined systems occur without explicit forensic authorization.

The solution enforces write-blocking through several mechanisms. Automount services are disabled to prevent automatic drive mounting, and the `blockdev --setro` command enforces read-only states at the block device level. When analysts need to save samples to external devices during analysis, a dedicated application enables controlled write access exclusively to external devices, with a visual indicator (a red icon) signaling when write access is active. This mechanism does not apply to the USB boot partition, thereby maintaining complete analysis segregation.

To address memory constraints on legacy hardware, the system uses high-throughput USB 3.0 external storage for swap and temporary file operations. The bootable ISO is deployed on a drive with at least 8 GB of capacity, with the remaining space forming an exFAT partition labeled IPED-TRIAGE. A startup script automatically allocates a swap file on this volume, offloading volatile data to prevent Out-of-Memory (OOM) failures on systems with limited RAM (e.g., 4 GB). The USB 3.0 interface provides the necessary I/O bandwidth for efficient virtual memory paging, avoiding the latency bottlenecks of USB 2.0.

Detection and handling of BitLocker-encrypted volumes are automatic, ensuring complete write-blocking. When BitLocker is in a suspended state, access to the protected volume is direct; otherwise, the 48-digit BitLocker recovery key or recovery password can be used to decrypt it, if provided. Windows Logical Volume Manager (LVM) partitions and Virtual Shadow Copies (VSS) are also supported, allowing recovery of deleted files present in backup snapshots but not in the current system state.

5.2 System Integration and Hybrid Architecture

A key aspect of the system is the union of hash-based and AI-based detection. Hash database filtering (Stage 1) enables high-speed execution (thousands of hashes per second) and avoids the reprocessing of known and duplicate files. AI resources focus on unknown content, and baseline intelligence comes from hash matches. AI classification (Stage 2) optimizes computational load for unknown files, achieving 95.8% CSAM recall on unknown content, directly addressing the limitations of traditional hash-based techniques and maintaining forensic interpretability through IPED integration.

Author Proof

The forensic triage workflow proceeds through evidence acquisition (forensically sound image from seized device) or live image initialization, case creation (IPED processes material through standard parsers), hash filtering (IPED applies ProjectVIC and NIST CAID sets to pre-filter known files), automated AI triage (CSAM detection plugin evaluates unknown images and video frames), prioritization (creating bookmarks for hash hits and categorizing AI detections according to the confidence level), manual verification (analyst reviews flagged content), and report generation (forensic documentation with maintained chain of custody).

6 Results and Performance Analysis

The training dynamics are in Fig. 1, which shows the evolution of training and validation accuracy and loss over 100 epochs (50 epochs of feature extraction and 50 of fine-tuning; not completed due to early stopping). The curves show stable convergence without significant overfitting, with validation metrics closely tracking training metrics throughout both training phases.

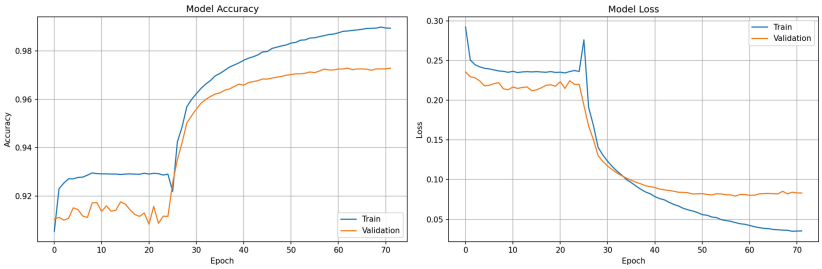

Fig. 1. Training history showing accuracy and loss evolution over 100 epochs (with early stopping). The two-stage training approach demonstrates stable convergence with minimal overfitting.

Table 2. Test Set Classification Report

Class	Precision	Recall	F1-Score	Support
csam	0.914	0.958	0.935	4,633
porn	0.985	0.970	0.977	13,083
other	0.998	0.996	0.997	24,215
Macro Avg	**0.965**	**0.975**	**0.970**	**41,931**

Evaluation on a realistic class distribution (11.0% csam, 31.27% porn, 57.8% other) across 41,931 held-out test samples yielded an overall accuracy of 96.5%. The detailed classification report is in Table 2. A 95.8% CSAM recall ensures

Author Proof

detection of nearly all unknown illicit material in field scenarios, while 91.4% precision results in acceptable false-positive rates for forensic triage, where human verification follows automated detection. Near-perfect performance on the other class (99.8% precision, 99.6% recall) demonstrates strong discrimination between innocuous and illicit content.

The confusion matrix (Table 3) illustrates the model's performance on the test set. As expected, there is a slight misclassification between the visually similar CSAM and pornography classes, while the 'other' class shows negligible error rates.

Table 3. Confusion matrix for the test set

		Predicted		
		CSAM	Porn	Other
True	CSAM	**4,440**	167	26
	Porn	357	**12,696**	30
	Other	62	31	**24,122**

Table 4. Inference Performance on Various Hardware

System	GPU	CPU	Img/s	ms/img
GPU Lab	NVIDIA 5700	Xeon Silver 4214	250	4
Moderate	Quadro K2000	Xeon E5-2630	48.1	20.82
Laptop	None	AMD 5300U	55	18
Legacy Laptop	None	Intel i5-7200U	35	28.5

End-to-end inference performance, from image loading to classification output, was evaluated across diverse hardware configurations, as detailed in Table 4. The results confirm the system's viability for field deployment: processing a typical seized device containing 100,000 images would require less than 60 min on field hardware, enabling same-day triage. The system supports both still-image and video content analysis by extracting frames from video files, enabling comprehensive multimedia analysis.

7 Discussion

The system achieved 91.4% CSAM accuracy on the test dataset, exceeding the 90% target. The hybrid detection workflow effectively integrates hash prefiltering with deep learning classification. Advanced preprocessing techniques prevent contamination between the training and test sets, and IPED integration

Author Proof

enables seamless integration into the workflow. The portable forensic environment supports both field and laboratory deployment, with performance validated across GPU, modern, and legacy CPU configurations.

This approach enables rapid triage during field operations, achieving 95.8% recall for unknown CSAM and directly addressing investigative gaps. The system processes 100,000 images in under one hour, facilitating same-day intelligence, and reduces unnecessary AI computation on already-identified files through hash analysis. The system also enables occupational health improvements for analysts through automated filtering, minimizing exposure to traumatic content, reducing occupational stress and psychological impacts, enhancing retention of specialized personnel, and supporting staff duty of care. Notably, possession of CSAM constitutes a felony in most jurisdictions. By accelerating the identification of illicit content, the system expedites investigations and may facilitate in flagrante delicto arrests.

Limitations of this study include training restricted to 46,662 CSAM images from the Brazilian Federal Police databases, which may introduce geographic and temporal bias. Technical constraints involve video processing via frame extraction, uncertain performance on heavily compressed or altered images, and limited support for transformation attempts. Operational constraints include the need to modify BIOS settings to enable the Linux environment on some systems, the lack of Secure Boot support, and the need for additional validation to ensure legal admissibility in certain jurisdictions.

8 Conclusions and Future Work

This study developed and validated a practical, semi-automated system for CSAM detection in resource-constrained forensic environments. By integrating transfer learning, lightweight CNN architectures, and hash-based pre-filtering within established forensic frameworks and portable Linux distributions, the proposed solution addresses critical operational, occupational health, and investigative challenges in digital forensics.

When deployed at scale, the system has the potential to transform investigative workflows by reducing case processing times from months to weeks and alleviating laboratory backlogs through rapid filtering. Beyond operational efficiency, the solution improves investigative outcomes by enabling consistent field triage and accelerating victim identification through the detection of unknown materials. The methodology also prioritizes occupational health by minimizing analyst exposure to traumatic content, and its replicability facilitates broader international investigations.

Future research will focus on enhancing technical robustness, forensic utility, and alignment with global standards. Planned developments include extending the video architecture with temporal CNNs and creating models for AI-generated CSAM and adversarial resistance. Additionally, leveraging the flexibility of the IPED tool, a larger model based on the same dataset is planned for release to support deployment on forensic workstations with GPUs.

Author Proof

Acknowledgments. The authors gratefully acknowledge the Brazilian Federal Police for providing access to a database of known illegal content under strict ethical protocols. Special thanks to the Department of Electrical Engineering at the University of Brasília for support. This work was conducted in strict compliance with Brazilian legal frameworks, including the LGPD, and ethical research standards.

References

1. Brazil: Law no 13.709, of August 14, 2018: General personal data protection law (lgpd) (2018). https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/L13709.htm. Accessed 2025
2. Cantrell, G., Dampier, D.A.: Implementing the automated phases of the partially-automated digital triage process model. J. Digit. Forensics Secur. Law **7**(1), 1–25 (2012). https://doi.org/10.15394/jdfsl.2012.1135
3. Casey, E.: Digital Evidence and Computer Crime: Forensic Science, Computers, and the Internet. Forensics 2011. Elsevier Science (2011)
4. Coltel, R.: Dislocker: Fuse driver to read/write windows' bitlocker-ed volumes under linux / mac osx (2024). https://github.com/Aorimn/dislocker. Accessed 08 Dec 2025
5. Danielle Crimmins Jonathan McGrath, J.B.: Safeguarding Forensic Science Professionals. IACP Publication (2024)
6. Du, X., et al.: Sok: exploring the state of the art and the future potential of artificial intelligence in digital forensic investigation. CoRR **abs/2012.01987** (2020). https://arxiv.org/abs/2012.01987
7. Gogia, G., Rughani, P.H.: ML based digital forensics framework. J. Digit. Forensics, Secur. Law **17**(6) (2023)
8. Gutfeter, W., Gajewska, J., Pacut, A.: Detecting sexually explicit content in the context of the child sexual abuse materials (csam): end-to-end classifiers and region-based networks (2024). https://arxiv.org/abs/2406.14131
9. He, H., Garcia, E.: Learning from imbalanced data. IEEE Trans. Knowl. Data Eng. **21**(9), 1263–1284 (2009)
10. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580 (2012). https://arxiv.org/abs/1207.0580
11. Hoang, D.A., Nguyen, C., Vasileios, B., Do, T.T., Carneiro, G.: Maximising the utility of validation sets for imbalanced noisy-label meta-learning (2025)
12. IPED Development Team: IPED: Indexador e Processador de Evidências Digitais (2024). https://github.com/sepinf-inc/IPED
13. Jajal, P., et al.: Analysis of failures and risks in deep learning model converters: a case study in the onnx ecosystem (2024). https://arxiv.org/abs/2303.17708
14. Jusas, V., Birvinskas, D., Gahramanov, E.: Methods and tools of digital triage in forensic context: survey and future directions. Symmetry **9**(4), 49 (2017). https://doi.org/10.3390/sym9040049
15. Kornblith, S., Shlens, J., Le, Q.: Do better ImageNet models transfer better? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2661–2671 (2018)
16. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)

Author Proof

17. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proc. IEEE **86**(11), 2278–2324 (1998). https://doi.org/10.1109/5.726791
18. Lee, H., Ermakova, T., Ververis, V., Fabian, B.: Detecting child sexual abuse material: a comprehensive survey. Forensic Sci. Int. Digit. Invest. **34**, 301022 (2020). https://doi.org/10.1016/j.fsidi.2020.301022
19. National Center for Missing & Exploited Children: CyberTipline Reports (2024). https://www.ncmec.org/
20. Pearson, S., Watson, R.: Digital Triage Forensics: Processing the Digital Crime Scene. Syngress (2010)
21. Rogers, M.K., Goldman, J., Mislan, R., Wedge, T., Debrota, S.: Computer forensics field triage process model. J. Digit. Forensics, Secur. Law **1**(2), 27–40 (2006). https://doi.org/10.15394/jdfsl.2006.1004
22. Roussev, V., Quates, C., Martell, R.: Real-time digital forensics and triage. Digit. Investig. **10**(2), 158–167 (2013). https://doi.org/10.1016/j.diin.2013.02.001
23. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.: MobileNetV2: inverted residuals and linear bottlenecks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4510–4520 (2018)
24. Scientific Working Group on Digital Evidence: SWGDE Best Practices for Computer Forensic Acquisitions. Technical report. SWGDE (2018)
25. Shavers, B.: Placing the suspect behind the keyboard. Syngress (2013). https://doi.org/10.1016/B978-1-59-749985-9.00001-0
26. Shaw, A., Browne, A.: A practical and robust approach to coping with large volumes of data submitted for digital forensic examination. Digit. Investig. **10**(2), 116–128 (2013). https://doi.org/10.1016/j.diin.2013.04.003
27. Shorten, C., Khoshgoftaar, T.: A survey on image data augmentation for deep learning. J. Big Data **6**(1), 60 (2019)
28. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. In: International Conference on Machine Learning, pp. 10,096–10,106 (2021). https://arxiv.org/abs/2104.00298
29. Weiss, K., Khoshgoftaar, T., Wang, D.: A survey of transfer learning. J. Big Data **3** (2016). https://doi.org/10.1186/s40537-016-0043-6
30. Wolbers, H., Cubitt, T., Cahill, M.: Artificial intelligence and child sexual abuse: a rapid evidence assessment. Trends Issues Crime Crim. Justice, **711** (2025). https://doi.org/10.52922/ti77802

Author Proof

Author Queries

Chapter 23

Query Refs.	Details Required	Author's response
AQ1	This is to inform you that corresponding author has been identified as per the information available in the Copyright form.	
AQ2	Please check and confirm if the authors Given and Family names have been correctly identified.	

Alternative Texts for Your Images, Please Check and Correct them if Required

Page no	Fig/Photo	Thumbnail	Alt-text Description
6	Fig1	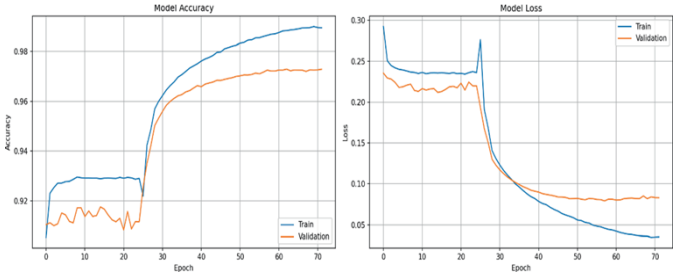	Two line charts compare training and validation results of a model over 75 epochs. The left chart shows accuracy, with the blue line for training and the orange line for validation. Both start around 0.92 accuracy, then after epoch 25, accuracy rises sharply, reaching about 0.99 for training and 0.975 for validation by the end. The right chart shows loss, with the same color scheme. Loss starts near 0.3 for training and 0.25 for validation, then drops sharply after epoch 25. Training loss decreases steadily to about 0.03, while validation loss levels off around 0.08. The charts highlight a key point at epoch 25 where model performance improves significantly, showing better accuracy and lower loss for both training and validation, indicating effective learning and generalization.