

Compliance Challenges in Forensic Image Analysis Under the Artificial Intelligence Act

Benedikt Lorch*, Nicole Scheler†, Christian Riess*

*IT Security Infrastructures Lab, Friedrich-Alexander-Universität, Erlangen, Germany

†International Criminal Law Research Unit, Friedrich-Alexander-Universität, Erlangen, Germany
{benedikt.lorch, nicole.scheler, christian.riess}@fau.de

Abstract—In many applications of forensic image analysis, state-of-the-art results are nowadays achieved with machine learning methods. However, concerns about their reliability and opacity raise the question whether such methods can be used in criminal investigations. So far, this question of legal compliance has hardly been discussed, also because legal regulations for machine learning methods were not defined explicitly.

To this end, the European Commission recently proposed the artificial intelligence (AI) act, a regulatory framework for the trustworthy use of AI. Under the draft AI act, high-risk AI systems for use in law enforcement are permitted but subject to compliance with mandatory requirements.

In this paper, we review why the use of machine learning in forensic image analysis is classified as high-risk. We then summarize the mandatory requirements for high-risk AI systems and discuss these requirements in light of two forensic applications, license plate recognition and deep fake detection. The goal of this paper is to raise awareness of the upcoming legal requirements and to point out avenues for future research.

Index Terms—forensic image analysis, artificial intelligence act

I. INTRODUCTION

Forensic image analysis subsumes a range of techniques to analyze the contents of an image or the image itself in legal matters [1]. Relevant sub-disciplines include photographic comparison, photogrammetry, source identification, and image authentication. The initial motivation for the development of these fields was to support criminal investigations by providing scientifically grounded statements about probative image material. Although forensic image analysis has now attracted the interest of a wider audience such as journalistic fact-checkers and insurance companies, the analysis of evidence in criminal investigations is still one of the field’s core interests.

After the recent success of deep learning, the research focus in multimedia forensics shifted from analytical models to machine learning. Machine learning is attractive for forensic image analysis for several reasons. First, advances in object detection, image denoising, and body pose estimation oftentimes directly translate to forensic tasks like license plate recognition [2], noise residual extraction [3], and photogrammetric height estimation [4]. Second, many forensic traces are notoriously difficult to model analytically, due to limited knowledge about the image origin (*e.g.*, camera hardware or

processing) and due to the complex interaction between traces and image transformations. In these cases, learning from large sets of examples provides an appealing solution. Additionally, machine learning methods excel in broader application scenarios, whereas the traditional analytical methods are usually confined to a narrow application scope [5].

Despite their high performance, machine learning methods suffer from a number of intrinsic shortcomings. Most importantly, their opacity makes it difficult to explain how the method works internally and to trace a particular output. Hence, an open challenge is how to audit and verify that a machine learning system works as intended and has no adverse side effects or biases. Additionally, machine learning methods are sensitive to training-test mismatches. This issue is particularly relevant in forensic applications, where methods are regularly exposed to images from uncontrolled origins, which may not be covered by the training data. In such cases, machine learning methods are prone to making wrong predictions [6]. Especially in applications where methods rely on imperceptible traces, detecting such mismatches is difficult.

These limitations raise the question to what extent machine learning tools for forensic image analysis can be applied to criminal prosecution cases, in all stages from the investigation phase to providing expert witness testimony in the courtroom. Although the use of clues from forensic image analysis in the investigation phase may appear less delicate, such clues can lead to further investigative measures that affect fundamental rights of the suspect. Investigative measures that restrict the fundamental rights of an individual must be based on a sufficient and secured factual basis. Therefore, forensic analysis methods must be as reliable as possible, and their legal compliance cannot be taken for granted.

Although current research focuses on learning-based methods, their practical use in law enforcement and their legal compliance are rarely discussed. This may be because legal compliance regulations either vary between jurisdictions, are unclear to the technical community, or even do not exist yet.

In April 2021, the European Commission proposed a regulatory framework for the trustworthy use of AI, the *artificial intelligence act* (AIA) [7].¹ Although the proposal is still

¹The AIA’s definition of AI technologies includes machine learning and deep learning, but also knowledge-based reasoning, expert systems, statistical approaches, and search and optimization methods. The Commission deliberately chose this broad definition of AI to include future developments [8].

Work was supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Research and Training Group 2475 ”Cybercrime and Forensic Computing” (grant #393541319/GRK2475/1-2019).

under discussion, the AIA will ultimately regulate the use of AI methods across the EU, including in law enforcement. Therefore, the AIA establishes legal clarity on which uses of AI are prohibited, which uses are permitted subject to mandatory safeguards, and which uses are permitted with few or no restrictions.

In this paper, we examine the AIA and analyze its impact on learning-based forensic image analysis methods. The AIA classifies forensic image analysis applications in law enforcement as high-risk. Under the AIA, AI systems for these applications are permitted but they must meet mandatory requirements. As our contribution, we summarize and discuss the most relevant requirements in light of two forensic applications, namely license plate recognition and deep fake detection. The goal of this paper is to raise awareness within the research community about legal requirements for AI tools, and to point out directions for future research motivated by the AIA.

For our discussion we choose these two applications, because machine learning is a promising approach in both cases, far outperforming traditional methods and humans. While license plate recognition has immediate application in law enforcement, deep fake detection is still a young field, but nevertheless highlights some of the key challenges towards making learning-based tools legally compliant under the AIA.

This paper is organized as follows: Section II reviews related work on legal requirements in forensic image analysis. Section III summarizes the AIA and the requirements to high-risk AI systems. In Sec. IV, we discuss these requirements using the two forensic applications of license plate recognition and deep fake detection. Section V concludes the paper.

II. RELATED WORK

Only few works discuss the applicability of forensic algorithms from a legal perspective. As a notable example, Kraetzer and Dittmann summarize the U.S. Federal Rules of Evidence and the Daubert criteria and discuss the maturity of camera device forensics (PRNU) in light of these compliance rules [5]. In a similar spirit, De Rosa *et al.* summarize ISO standards and guidelines related to digital investigations and discuss application scenarios of traditional media forensics methods [9]. However, neither these ISO standards nor the Daubert criteria give concrete advice about AI methods. Moreover, the Daubert criteria only apply to expert witness testimony in the U.S., but not during the investigation phase.

Furthermore, the Scientific Working Group on Digital Evidence (SWGDE) and the European Network of Forensic Science Institutes (ENFSI) released a series of guides to promote best practices and to establish comparability in forensic casework. These documents cover, for example, image authentication, forensic image comparison, and requirements for testing forensic tools. However, these documents do not yet consider the use of AI, with one exception: In [10], the SWGDE outlines the emerging use cases in law enforcement and the challenges related to AI-based video analysis. While acknowledging the potential benefits of AI tools, the guide

cautions investigators to be aware of the limitations of AI and to diligently follow established standards. Similar to best practice guides, forensic process models guide investigations and facilitate comparability and reproducibility [11]. However, both these guides and process models do not yet take into account the particular challenges related to AI methods.

In the related field of biometrics, the use of AI for facial recognition has raised several concerns, including the risk of discrimination, privacy aspects of large facial databases, and the potential use for public surveillance [12]. Because these concerns outweigh the benefits, the AIA prohibits facial recognition systems for public surveillance, with special exceptions for law enforcement. In this paper, however, we focus on other applications of forensic image analysis such as license plate recognition and deep fake detection. For a detailed discussion of facial recognition, we refer to the in-depth analysis by the European Parliament Research Service [13].

III. ARTIFICIAL INTELLIGENCE ACT

This section first provides general background on the AIA. Then we discuss why forensic image analysis applications used by law enforcement are classified as high-risk. Lastly, we summarize the requirements to high-risk AI systems.

A. Background

In the AIA, the European Commission proposes a horizontal legal framework for the trustworthy use of AI [7]. This proposal is based on recommendations by an appointed high-level expert group [14], a preceding white paper, and public stakeholder consultation. The AIA applies to both private, public, as well as extraterritorial providers whose AI tools are used in the EU. Once adopted (presumably in 2024 at the earliest [15]), the AIA will become directly applicable law in all EU member states.

The AIA classifies the use of AI in four risk categories called unacceptable risk, high risk, transparency risk, and minimum risk. Applications with unacceptable risk are prohibited. High-risk AI systems are permitted but subject to mandatory requirements. Applications with transparency risks or minimum risks are permitted subject to transparency obligations or without restrictions, respectively. This paper focuses on the high-risk category, in which we see many forensic image analysis techniques when used by law enforcement.

The AIA leaves the technical implementation of these requirements to the providers of the respective AI systems, so that providers can take into account the scientific and technical progress. Nevertheless, the Commission will instruct the European standardization organizations to translate these requirements into harmonized technical standards. If providers adhere to these standards, their AI systems are given a presumption of conformity (Art. 40). These harmonized standards will further clarify and technically define some of the still abstract requirements stated in the AIA. Nevertheless, we believe that the research community benefits from an early discussion of these requirements, such that these requirements can be taken into consideration in future research.

B. High-risk AI systems

Annex III explicitly states law enforcement application scenarios that the Commission classifies as high-risk. These include AI systems for evaluating the reliability of evidence in the course of investigation or prosecution of criminal offenses, and AI systems for deep fake detection. Therefore, we conclude that AI-based image authentication counts as high-risk application.

It is not explicitly stated whether AI systems for source identification or photographic comparison, including license plate recognition, are considered high-risk. However, the Commission is entitled to expand the list of high-risk applications. AI systems can be added to the list when they pose a similar or even higher risk to the health, safety, or fundamental rights than the applications already listed. Since source identification and photographic comparison aim at identifying a suspect, potentially leading to investigative measures against the suspect, we expect these applications to also be considered high-risk.

C. Requirements to high-risk AI systems

We now summarize requirements to high-risk AI systems. We focus on requirements that are relevant to the development of the AI model, leaving aside operational requirements.

1) *Accuracy, robustness and cybersecurity (Art. 15)*: High-risk AI systems must achieve an appropriate level of accuracy, robustness, and cybersecurity for their intended purpose. AI systems must tolerate errors, faults and inconsistencies that may appear within their working environment. AI systems must be resilient to attacks, in particular related to manipulations of the training dataset (data poisoning), misleading inputs (adversarial examples), and model flaws.

2) *Risk management and testing (Art. 9)*: The known and foreseeable risks associated to the AI system should be identified, evaluated, and documented systematically within a risk management system. To alleviate these risks, mitigation strategies should be implemented. High-risk AI systems must be tested to ensure they perform consistently for their intended purpose and to validate the risk mitigation strategies.

3) *Data and data governance (Art. 10)*: Training, validation, and testing data are subject to appropriate data governance practices. These practices concern the data collection and pre-processing, data assumptions, data availability and quantity, and the examination of possible biases, among other things. In particular, all data sets must be relevant, representative, free of error, and complete. Datasets should take into account the characteristics of the individuals and the geographical, behavioral, or functional settings for which the AI system will be used.

4) *Transparency (Art. 13) and human oversight (Art. 14)*: The AI system must be sufficiently transparent such that users can interpret and use the AI system's output appropriately. The degree of transparency must enable the user to oversee the AI system in order to prevent and minimize risks to health, safety, and fundamental rights. Providers must specify measures that allow a human operator to fully understand the capabilities and limitations of the AI system, to correctly

interpret the AI system's output with the tools available, to decide to disregard the AI system's output, and to remain aware of human tendency to over-rely on the AI system.

5) *Documentation*: The AI system must be accompanied by several sources of documentation including an instruction manual (Art. 13), a technical documentation (Annex IV), a risk management system (Art. 9), and a quality management system (Art. 17). The instruction manual specifies the AI system's capabilities and limitations. This includes its intended purpose, the expected level of accuracy, robustness and cybersecurity, known and foreseeable circumstances that impact these performance metrics, specifications for the input data, and the interpretability measures. The technical documentation describes in detail the system components and its development including conceptual decisions. The risk management system describes the risks associated with the use of the system and how they are mitigated. The quality management system documents compliance with all regulations of the AIA.

IV. COMPLIANCE CHALLENGES

In this section, we review the requirements from Sec. III-C and discuss to what extent these requirements are currently addressed in research. We specifically use the examples of license plate recognition and deep fake detection.

1) *Accuracy, robustness, cybersecurity*: In simple terms, this requirement demands AI systems to perform well. Overall, we find that these requirements align with the goals and established performance metrics in the research community.

For example, recent work on license plate recognition aimed at increasing accuracy and improving robustness to image degradations compared to prior work [2]. We are not aware of studies on the adversarial robustness of forensic license plate recognition. However, the input images usually stem from security cameras with restricted access. Therefore, a feasible direction to prevent attacks could be to actively embed security schemes for image authentication. Nevertheless, a more common issue than tampering are image degradations due to environmental aspects such as rain, dirt, and difficult lighting conditions. Future work could study the robustness w.r.t. these degradations.

In deep fake detection, a recent line of work focuses on improving the accuracy and robustness in light of unseen manipulations. Several works also studied the resilience of deep fake detectors to evasion attacks [16]. Despite these developments, improving robustness to real-world image degradations and unknown manipulations as well as defending against attacks are still open problems, but are being actively researched.

2) *Risk management and testing*: The AIA requires providers to thoroughly test their AI systems. Because AI systems do not give theoretical guarantees, testing is essential to establish the required performance metrics, to identify risks, and to validate risk management strategies. So far, there are no established or standardized protocols for AI testing [17].

Although technical research papers usually include an experimental validation, these validations are often limited in breadth and depth, they only focus on a particular performance

metric, and testing protocols vary between papers. In contrast, testing AI systems for real-world applications requires a structured testing protocol. This testing protocol should assert, in addition to performance, quality criteria such as fairness, safety and reliability. Yet, this testing protocol must be sufficiently flexible to support different notions of these quality criteria, because the notions of fairness and safety vary between different applications.

To this end, the *ExamAI* project [17] proposed *assurance cases* as an argumentation framework to structurally document how an AI system meets the required quality characteristics. In a nutshell, assurance cases consist of a normative claim, arguments to support the claim, and factual evidence to support these arguments. Assurance cases therefore make transparent to both developers and external auditors what has been experimentally validated and how this factual evidence supports the claim of these qualities. As long as there are no standardized protocols, we see assurance cases as a helpful tool.

However, we note that large-scale testing is likely beyond the capacity of individual researchers and requires assistance from companies, research institutes, or federal agencies. Nevertheless, the required risk management system provides motivation to research structured testing protocols (*e.g.*, assurance cases for license plate recognition and deep fake detection), to identify risks associated to forensic AI applications, and to develop mitigation strategies. For example, future work on detecting deep fakes could investigate the risk of discrimination, where first results showed that detectors tend to misclassify people with skin diseases and overly smooth faces [18].

3) *Data and data governance*: The AIA prescribes training, validation, and testing datasets to be relevant, representative, free of error, complete, and datasets must be examined for biases. Arguably, the meaning of these properties must be discussed in the context of each application.

In the context of license plate recognition, relevance could mean that all images must show license plates, and error-free means that labels are correct. Representativity and completeness can be interpreted such that images must cover all types of image degradations and geographical variations relevant for the actual use case. A typical source of bias could be an over- or under-representation of some characters or a geographic region. Unfortunately, publicly available license plate datasets are small, which limits their representativeness and makes them too incomplete to reflect real-world conditions. For this reason, recent works trained their AI models with synthetic images [2], [19]. While synthetic training data mitigates some of these issues, synthetic images must be representative enough such that the model generalizes to real images.

The current trend towards larger models creates demand for even larger datasets. Because data collection is very expensive, compiling a complete and representative dataset is, again, likely beyond the capacity of individual researchers. Although scraping images from the internet is an appealing shortcut, such datasets are likely to contain biases or labeling errors.

We welcome large datasets provided by governmental organizations, *e.g.*, in the DARPA MediFor program, or private

organizations, *e.g.*, Facebook’s deepfake detection challenge dataset. This latter dataset features a wide variety of gender, skin, tone, ethnicity, age, and other characteristics, and it contains a large number of image degradations. Nevertheless, these datasets can still contain biases and their large size makes them difficult to screen.

An interesting direction for future research could therefore be to develop automatic screening tools and metrics to assess the quality of datasets. As an example, recent work identified numerous labeling errors in several popular datasets including ImageNet through an automated method [20]. Another open problem is how to document that datasets meet the required properties. Future work could develop assurance cases to assert these data quality properties in a structured manner.

4) *Transparency and human oversight*: Interestingly, the AIA in its current draft requires only sufficient transparency, while the high-level expert group called for a stronger notion of explainability, namely that humans can understand and trace predictions by an AI system [14]. This weakened requirement leaves opportunities for the use of black-box models such as neural networks. Nevertheless, the degree of transparency achieved by the AI system must enable humans to correctly interpret and potentially disregard the AI system’s output. To our understanding, this ability to disregard the AI system’s output requires a way to identify malfunctioning. To this end, we see two complementary strategies: explaining model predictions and uncertainty-aware detectors.

Work on *explainable AI* (XAI) can be categorized into inherently interpretable models and post-hoc explanations. Two inherently interpretable approaches for deep fake detection rely on hand-crafted features and use a random forest [11] or kNN classifier [21]. These methods perform well when the forensically relevant traces are strong, as it is the case in constrained scenarios. But when forensically relevant traces are weak, *e.g.*, statistical traces after resizing and compression, deep learning outperforms hand-crafted features [16]. For deep learning models, a few post-hoc explanation methods have been applied to deep fake detectors [22], but so far only studied in very limited settings. In license plate recognition, first works studied saliency maps [23] and a sequential architecture with intermediate inspection possibilities [19].

However, we share the opinion that current post-hoc XAI approaches are incapable of providing reliable explanations for every individual prediction [24], [25]. This is because simple approximations may not be faithful to the original black-box model and because human interpretations based on post-hoc explanations may diverge from the actual model computation. Even though post-hoc XAI cannot provide explanations for every prediction, these visualizations can nevertheless indicate failure cases. For example, a deep fake detection result can be deemed unreliable if a saliency map points at the background region even though the detection should focus on the foreground. Therefore, post-hoc XAI approaches can at least unveil some cases of malfunctioning. Given that the AIA does not specify the required level of interpretability, it is worth exploring both directions: stronger inherently interpretable

models and post-hoc explanations. We note that explaining AI methods for forensic tasks may be more difficult than for image recognition, because the forensically relevant traces are often imperceptible. Therefore, an open challenge is to visualize the meaning of these traces even to experts.

A complementary approach to identifying erroneous predictions could be uncertainty-aware methods, *e.g.*, through Bayesian modeling [6]. These methods provide an uncertainty with every prediction. This enables the operator to quantify how much to trust in the prediction. So far, we are not aware of uncertainty-aware detectors for license plate recognition or deep fake detection. Despite first results, the reliability of these detectors has not been studied on a large scale.

5) *Documentation*: In research, technical papers sometimes fall short of discussing limitations of a proposed method, assumptions on the input data, and potential risks associated to an AI application. We see this documentation requirement as a motivation to increase transparency also in research papers.

In fact, major conferences such as *NeurIPS 2021* introduced a paper checklist for responsible machine learning research.² Interestingly, this checklist overlaps with the AIA documentation requirements in terms of discussing limitations and assumptions, identifying potential risks with negative societal impact, and requiring necessary technical details for reproducibility. At the same time, reviewers are instructed to acknowledge authors for discussing limitations and societal implications. Similarly, IEEE T-IFS published guidelines for deep learning submissions³, although these only focus on reproducibility, leaving aside data governance or risks.

After all, improving transparency in research papers would not only foster reproducibility. It would also allow the research community to understand potential risks, to develop mitigation strategies, to establish best practices in data governance, and to understand limitations of state-of-the-art methods.

V. CONCLUSION

Recently proposed methods in forensic image analysis are increasingly based on machine learning. Despite being motivated by law enforcement applications, the practical use of AI tools and their legal compliance are rarely discussed. To this end, the AIA establishes legal clarity that AI systems in such high-risk applications are permitted but subject to mandatory requirements. In this paper, we reviewed these requirements and discussed their alignment with currently acknowledged goals of the research community. Although it is the providers of AI systems who need to show compliance with the AIA's regulations, some of these requirements present interesting directions for future research, *e.g.*, identifying risks and mitigation strategies, developing data governance practices, and improving interpretability. While the AIA is still subject to amendment, we hope this paper will stimulate discussion about the practical applicability of AI tools in multimedia forensics.

²<https://neuripsconf.medium.com/introducing-the-neurips-2021-paper-checklist-3220d6df500b>, accessed on 16.02.2022

³https://signalprocessingsociety.org/sites/default/files/uploads/publications_resources/docs/Guidelines_deep_learning_submissions.pdf, accessed on 16.02.2022

REFERENCES

- [1] Scientific Working Group on Digital Evidence, "SWGDE guidelines for forensic image analysis," <https://drive.google.com/file/d/1y549OzGFK9QphnVOVtj1W2yUIIjSroF5>, Feb. 2017.
- [2] D. Moussa, A. Maier, F. Schirmacher, and C. Riess, "Sequence-based recognition of license plates with severe out-of-distribution degradations," in *Comp. Anal. Images Patterns*, 2021, pp. 175–185.
- [3] D. Cozzolino and L. Verdoliva, "Noiseprint: A CNN-based camera model fingerprint," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 144–159, 2020.
- [4] N. Thakkar and H. Farid, "On the feasibility of 3D model-based forensic height and weight estimation," in *IEEE/CVF Conf. Comput. Vis. Pattern Recog. Workshops*, Jun. 2021, pp. 953–961.
- [5] C. Kraetzer and J. Dittmann, "Considerations on the benchmarking of media forensics," in *Eur. Signal Process. Conf.*, 2015, pp. 61–65.
- [6] B. Lorch, A. Maier, and C. Riess, "Reliable JPEG forensics via model uncertainty," in *IEEE Int. Workshop Inf. Forensics Security*, 2020.
- [7] European Commission, "Proposal for a regulation of the European parliament and of the Council laying down harmonised rules on AI (artificial intelligence act) and amending certain union legislative acts," Apr. 2021. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- [8] G. Spindler, "Der Vorschlag der EU-Kommission für eine Verordnung zur Regulierung der KI - Ansatz, Instrumente, Qualität und Kontext," *Computer und Recht*, vol. 37, no. 6, pp. 361–374, 2021.
- [9] A. De Rosa, A. Piva, M. Fontani, and M. Iuliani, "Investigating multimedia contents," in *Int. Carnahan Conf. Security Tech.*, 2014.
- [10] Scientific Working Group on Digital Evidence, "SWGDE overview AI trends in video analysis," https://drive.google.com/file/d/IEQuM_hpQxo0AwvceP6zZyeNroIp2fOc, Jan. 2021.
- [11] D. Siegel, C. Kraetzer, S. Seidlitz, and J. Dittmann, "Media forensics considerations on deepfake detection with hand-crafted features," *J. Imag.*, vol. 7, no. 7, 2021.
- [12] C. Rückert, "Mit KI auf Verbrecherjagd," *Verfassungsblog*, 2021.
- [13] T. Madiaga and H. Mildebrath, "Regulating facial recognition in the EU," Sep. 2021. [Online]. Available: [https://www.europarl.europa.eu/thinktank/en/document/EPRS_IDA\(2021\)698021](https://www.europarl.europa.eu/thinktank/en/document/EPRS_IDA(2021)698021)
- [14] European Commission and Directorate-General for Communications Networks, Content and Technology, *Ethics guidelines for trustworthy AI*. Publications Office, 2019.
- [15] M. Veale and F. Z. Borgesius, "Demystifying the draft EU artificial intelligence act - analysing the good, the bad, and the unclear elements of the proposed approach," *Computer Law Review Int.*, vol. 22, no. 4, pp. 97–112, 2021.
- [16] L. Verdoliva, "Media forensics and deepfakes: An overview," *IEEE J. Selected Topics Signal Process.*, vol. 14, no. 5, pp. 910–932, 2020.
- [17] Gesellschaft für Informatik, "Abschlussbericht ExamAI – KI Testing und Auditing," <https://testing-ai.gi.de/veroeffentlichungen>, 2021.
- [18] A. Prabhu, J. Materzyńska, P. K. Dokania, P. H. S. Torr, and S.-N. Lim, "Is your face fake? social impacts of algorithmic 'fake' determination," <https://fb.watch/b3dXk4QbAm/?t=2460>, Jun. 2020.
- [19] G. Rossi, M. Fontani, and S. Milani, "Neural network for denoising and reading degraded license plates," in *Int. Conf. Pattern Recog. Workshops*, 2021, pp. 484–499.
- [20] C. Northcutt, L. Jiang, and I. Chuang, "Confident learning: Estimating uncertainty in dataset labels," *J. Artif. Intell. Res.*, vol. 70, pp. 1373–1411, 2021.
- [21] F. Matern, C. Riess, and M. Stamminger, "Exploiting visual artifacts to expose deepfakes and face manipulations," in *IEEE Winter Appl. Comput. Vis. Workshops*, 2019, pp. 83–92.
- [22] B. Malolan, A. Parekh, and F. Kazi, "Explainable deep-fake detection using visual interpretability methods," in *Int. Conf. Inf. Comput. Tech.*, 2020, pp. 289–293.
- [23] B. Lorch, S. Agarwal, and H. Farid, "Forensic reconstruction of severely degraded license plates," in *Electronic Imaging*, Jan. 2019.
- [24] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Mach. Intell.*, vol. 1, no. 5, pp. 206–215, May 2019.
- [25] M. Ghassemi, L. Oakden-Rayner, and A. L. Beam, "The false hope of current approaches to explainable AI in health care," *The Lancet Digital Health*, vol. 3, no. 11, pp. 745–750, 2021.