Fake Invoices and Receipts in Accounting: Impacts and Mitigation in the Age of gen-AI

Philip Thornton[1,2]
P.Thornton2@newcastle.ac.uk
Lamyaa Aljuaid[1]
L.Z.M.Aljuaid2@newcastle.ac.uk
Matthew W Chapman[3]
M.Chapman10@newcastle.ac.uk
Deepayan Bhowmik[1]
Deepayan.Bhowmik@newcastle.ac.uk

[1] School of Computing, Newcastle University, Newcastle upon Tyne, NE4 5TG, United Kingdom.
[2] Sage PLC, Newcastle upon Tyne, NE28 9EJ, United Kingdom.
[3] School of Engineering, Newcastle University, Newcastle upon Tyne, NE1 7RU, United Kingdom.

Abstract

The advancement and accessibility of artificial intelligence (AI) present both opportunities and challenges for the accounting industry. This paper examines, specifically, the impacts of the threat posed by AI-generated fake invoices and receipts, explores the technology that underpins this and analyses the difficulties in their detection using traditional methods. The paper details the consequences for stakeholders, including accountants in practice, accountants in industry and small business owners in areas such as employee fraud, accounts payable fraud, tax compliance, audit reliability, forensic accounting, regulatory compliance, ethical considerations, and broader economic repercussions. Finally, it proposes a framework of mitigation, highlighting innovations such as AI-powered detection tools, e-invoicing, and data provenance, alongside process and regulatory improvements to safeguard the integrity and trust of accounting data in an increasingly AI-driven landscape. We propose a generative-AI-based manipulation pipeline and a detection tool, considered to be the first of its kind, and the result shows promising performance on standard invoice and receipt benchmark datasets.

1 Introduction

Trust forms the foundation of accounting workflows in which accuracy, transparency and reliability are critical for decision making and compliance. The accounting profession has a history of navigating technological shifts that challenge established processes, which demand new forms of trust, such as the transition to accounting software and the move to cloud-based technologies. The increasing integration of Artificial Intelligence (AI) in accounting workflows presents both opportunities and challenges [18]. The challenge of fraudulent documentation in invoice and receipt processing has long been recognised by accountants [4], with manual or semi-automated methods frequently used by fraudsters to manipulate documents using readily available tools. Established accounting practices currently incorporate various checks and controls, such as bank reconciliation and supplier verification, to mitigate these historical risks.

False invoicing occurs where a business fabricates or inflates invoices which name the business as the debtor. This allows it to fraudulently claim expenses for tax purposes that have not been incurred [21]. Historically, these fraudulent activities relied on manual methods and basic tools. For example, employees could submit fake receipts for non-existent purchases, created using basic image editing software, or by hand, or legitimate documents could be manually modified to inflate amounts, change dates, or adjust item descriptions, which involves cut-and-paste techniques, word processing or spreadsheet manipulation. Another common tactic is the creation of invoices from non-existent or shell companies to convince accounts payable teams into making unauthorised payments [20]. Additionally, fraudsters may exploit systems by submitting the same legitimate invoice multiple times for payment. These traditional methods, while they can be effective, are often limited by the fraudster's manual skill, the time required and the difficulties in producing large volumes of convincing fakes. Accounting practices have developed various checks and controls to mitigate these risks, which include manual inspection of documents, cross-referencing with physical supporting documentation such as purchase orders and emails, and the use of rule-based systems to flag anomalies based on set criteria, such as high invoice values, suspicious payment destinations or duplicate entries.

However, the advancement of AI introduces a different dimension to the challenge of financial fraud. Generative AI (Gen-AI) models are transforming the landscape of document forgery [16], making it easier to create sophisticated fake receipts and invoices. While some initial perspectives within the accounting community suggest that AI-generated or manipulated fake invoices do not present a new issue, this paper argues that the sophistication and scalability of AI-generated forgeries pose an increased threat that has the potential to overwhelm accounting processes, requiring a re-evaluation of current practices and the adoption of more suitable mitigations.

The impact of AI on the accounting industry has been the subject of increasing attention, with studies exploring AI's potential in automation and improved audit quality [17]. While these highlight the benefits of AI, the literature specifically addressing the threat of AI-generated fake financial documents and the implications for accountants is still growing. This paper aims to bridge this gap by providing an examination of the specific impacts of AI-generated fake financial documents on the accounting industry and extends beyond the technological capability to delve into the practical implications for accountants. This paper contributes by identifying the key threats and proposing a framework of mitigation that integrates technology with process improvements, addressing the limitations of traditional methods of fraud prevention. In supporting such a framework, this paper also proposes a first-of-its-kind pipeline for invoice and receipt manipulation by generative AI, a new dataset and a detection mechanism.

2 Impact Analysis

The use of easily accessible generative AI models, such as ChatGPT, has fundamentally altered the landscape of document forgery [23]. Previously, creating convincing fake receipts or invoices demanded a degree of specialised graphic design skills and considerable time investment; however, today's AI models can generate highly realistic accounting documents, in seconds, and at scale based on simple text prompts. While manually created fakes often contain visual inconsistencies that a trained eye can detect, AI can generate documents with a level of detail that surpasses human capabilities and can replicate subtle details found in

genuine documents, making reliance on visual checks alone unsuitable [28].

Gricius and Belovas [13] developed a tool to create realistic, fake invoices primarily for training machine learning models, as securing real invoice data is difficult due to privacy concerns and a lack of variety. Their method avoids advanced AI for generating the invoice content and layout. They deliberately chose this path due to common issues such as generating incorrect data, but acknowledge that techniques such as GANs and diffusion models could be used in future work. Tools such as ChatGPT can generate coherent, professional, and contextually appropriate messages which allow fraudsters to produce plausible email threads, impersonating suppliers or internal personnel, which align with fake invoices to commit fraud.

Fake invoices are easier to spot if they deviate from established patterns, such as invoice numbering. However, AI can be trained on large datasets to generate fake documents that align with historical data, such as supplier invoice numbering and item lines, potentially bypassing anomaly detection based manual rules. Wahid and Hassini [26] discussed the need for hybrid fraud detection frameworks, which look at the clustering of seemingly plausible invoice data for human review.

Accountants and SMEs may lack the time and resources to audit every transaction with the required scrutiny to detect an AI-generated fake, while AI's adaptability compounds the threat as adversarial training can take place creating an ongoing battle between fraud generation and detection. Okamoto *et al.* [22] further highlights this by investigating strategies for generating more challenging deep document forgeries to improve detection methods. Therefore, while accounting practices have created a foundation for fraud detection based on manually created fake documents, the unique capabilities of AI in generating realistic and scalable forgeries require a re-evaluation of these methods and the adoption of more advanced, technology-driven mitigations.

2.1 Generative AI in the accounting industry

The use of AI-generated fake receipts and invoices carries significant consequences for the accounting industry while eroding trust in digital workflows among accountants and Small and Medium-sized Enterprises (SME). One key area affected is employee fraud via expense management, where organisations face a heightened risk of fraudulent expense claims submitted by employees utilising AI-generated receipts for false expenditure or inflated amounts, with 24% of employees already admitting to expense fraud and a further 15% considering it [14]. While employees have always been able to submit fake receipts, AI allows for the creation of receipts that are more convincing and harder to audit, leading to a potential increase in fraudulent claims and financial losses, damaging employee/employer trust, and necessitating costly investigations. In addition, accounts payable (AP) fraud is heightened by AI's ability to generate realistic invoices from seemingly legitimate suppliers, coupled with AI-created emails that can bypass traditional verification processes. This vulnerability is particularly an issue for SMEs, who are likely to have fewer internal controls.

The impact also extends to tax fraud, as individuals and businesses may utilise AI to fraudulently reduce tax liabilities, leading to revenue losses for governments [7]. Accountants will be at the front line of this threat and will be required to review current audit processes to ensure they are set up to deal with these challenges, while forensic accountants will require new digital tools and techniques to identify AI-generated fake documents. The rise of such financial fraud also poses challenges for regulatory bodies with existing standards

such as FRS102 [12], which are built on the principle of presenting a true and fair view of a company's financial position.

AI's ability to generate fraudulent documents, raises ethical concerns, both for developers and potential misuses. The accounting profession must engage in discussions regarding the risks of these types of forgeries.

2.2 Economic and financial consequences

The use of AI-generated fake invoices and receipts poses a significant economic threat that goes beyond any immediate financial loss incurred by an organisation [11]. Companies face increased exposure to unauthorised outgoings via fake invoices paid to fraudsters' accounts and inflated or fictitious employee expenses (Expense Management), [14]. In its 2024 report, the Association of Certified Fraud Examiners highlighted losses, covering 1,921 cases and more than 3.1 billion USD. For larger entities, the volume of these fakes can lead to financial losses that are difficult to quantify. The US Government and the Government Accountability Office, 2024, estimate that federal programs lose between 233 billion USD and 521 billion USD to fraud annually [25].

Another significant economic consequence arises from the escalating costs of fraud prevention. As traditional controls prove inadequate, organisations will invest heavily in tools such as AI-powered fraud detection technologies, enhanced cybersecurity and employee training. The implementation and ongoing maintenance of these systems represent a significant expenditure [8]. Furthermore, when fraud is detected, the costs associated with forensic accounting investigations, legal fees, regulatory fines, and recovery efforts will bring additional costs and divert attention away from the running of the business. Additionally, this threat erodes trust and confidence within the broader economic ecosystem [4].

Undetected or highly publicized instances of AI-driven financial fraud can undermine investor confidence in companies, leading to decreased stock valuations or difficulties in securing investment. This erosion of trust can lead to increased due diligence costs, slower payment cycles, and an increased hesitancy in adopting digital ways of working if their security cannot be guaranteed.

3 Mitigation Strategies

Considering the limitations of traditional methods against this threat, mitigating the risks posed by AI-generated fake receipts and invoices requires a strategy that involves technological innovation blended with process improvement.

3.1 Remove reliance on paper-based documentation

Encouraging the widespread adoption of real-time digital receipt and invoice systems is a critical step in reducing the risks posed by AI-generated fraud. Paper, or image-based invoices, can be manipulated easily using genrative AI, wheras digital systems such as Europe's e-invoicing mandate, reduce this risk. Legislative initiatives such as Europe's e-invoicing mandate and secure transmission protocols such as PEPPOL (Pan-European Public Procurement Online) will assist in reducing this transition [10].

These systems introduce safeguards such as machine-readable formats, sender authentication and real-time reporting. This makes fraud harder to execute. However currently

adoption of these technologies is uneven. With small businesses and cross-border operations relying on manual invoices, in addition to this while e-invoicing employs comprehensive transmission security, internal threats still remain including threats such as compromised credentials or rogue employees, introducing fraudulent data into the e-invoicing system, meaning that human oversight still remains essential.

Finally, it can be determined that human vigilance remains key to stopping the threat of fraudulent activity.

3.2 Development of document integrity verification tools

The creation of tools, leveraging generative-AI techniques trained on large datasets containing both genuine and fake documents, will support in identifying anomalies, including inconsistencies in pixel patterns, metadata, and unnatural language patterns that may avoid human detection.

Alongside these detection tools, provenance technologies such as cryptographic hashing, digital signatures, and structured metadata such as JPEG Trust [6] can be embedded into documents to create a verifiable trail of their origin and history. By embedding provenance into the document lifecycle, recording who created it, how it has been modified, and who has been involved in this process, accountants can assess not just the integrity of the document's content but also the trustworthiness of its history.

Detection systems could be strengthened with secure AI pipline, which record provenance data at each stage[5]. This enables end-to-end accountability, which allows accountants to verify the integrity of the AI system.

3.3 Anomaly detection tools

Evolving to more sophisticated data analytics techniques, such as behavioural pattern recognition and predictive modelling, could help identify suspicious transactions and expense claims. AI can support with analysing typical expenditure, interactions and communication patterns with customers and suppliers to flag unusual activity that wouldn't be obvious through standard human-led processes. As highlighted by Schreyer *et al.* [24], an adversarial autoencoder can be used to spot unusual patterns in accounting data by understanding the normal structure of the data and flagging anything that deviates from it. Meanwhile, Wahid and Hassini [26] proposed a framework that combines different types of AI tools to group invoices into clusters based on their features (*e.g.*, invoice item lines). These clusters allow the system to highlight groups of invoices that don't follow usual patterns, ensuring the focus is on reviewing suspicious transactions more efficiently.

3.4 Build anti-fraud alliances

3.4.1 Cross-industry alliances, interoperability, standards and policy

Collaboration between accountants, technology developers, regulators, and law enforcement agencies, are vital for sharing intelligence. These alliances speed up knowledge. transfer, reveal common pattens and assist in the development of new standards and legislation. Global regulatory engagement is essential due to the nature of cross-border digital document exchange, and is supported by frameworks such as PEPPOL.

3.4.2 Continuing professional development and training

Accounting professionals at all levels need comprehensive education and training on the risks posed by AI-generated fake documents and the techniques for identifying fraud. Creating a new culture of an increased critical evaluation of documentation, informed by an understanding of AI capabilities, remains paramount [18]. This education and training must be embedded into the formal accounting qualifications process in collaboration with industry bodies and education providers.

3.5 Recommendations

To respond effectively to the growing threat of AI-generated fake financial documents, the accounting industry must pursue a collaborative, forward-looking strategy that covers all areas of the mitigation strategies identified above including: -

1. Removing reliance on paper-based documentation and looking at e-invoicing systems and legislation to support with this.
2. Development of document integrity verification tools to identify AI-generated manipulated invoices and receipts.
3. Embedding anomaly detection tools into accounting workflows to support with the analysis of large datasets and identify suspicious transactions.
4. Developing cross-industry alliances, interoperability, standards and policy to ensure a joined-up approach across the industry that shares best practice, identifies trends and champions relevant legislation.
5. Continuing professional development and training for accountants to ensure the role of the accountant is evolving with the threats associated with AI-generated fake invoices and receipts.

In support of the recommendations, this paper also proposes a pipeline for generating fake invoices and receipts using generative AI, supporting the development of a detection tool. The proposed technique (as described in Section 4) enables both the binary detection of the manipulated documents and localisation of the manipulated region.

4 Manipulation and Detection Pipeline

4.1 Fake invoice and receipt dataset generation

We developed an AI-manipulated invoice and receipt dataset using a semi-automated pipeline as depicted in Figure 1. First, the Optical Character Recognition (OCR) [19], a deep learning-based technique, is applied to detect the text within the input receipt. In order to differentiate between the extracted texts and facilitate manipulation, an index map is generated by assigning a numerical index to each extracted text in the list derived from OCR. Then, for controlling the text modification process and ensuring that the changes are related to the context of the invoice, such as seller names and amount, the prompt is manually typed by selecting the index of the target text from the table that contains the details derived from the OCR, such as extracted text, its confidence scores, and coordinates. Once the targeted

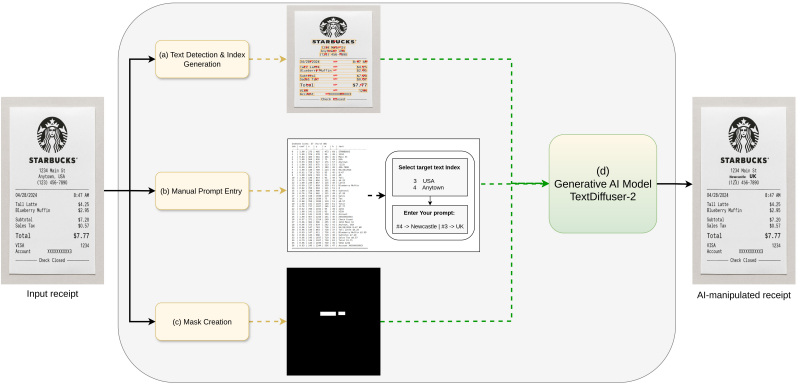

Figure 1: Overview of a semi-automated pipeline for generating AI-manipulated receipts and invoices. (a) Text detection & index generation, (b) Manual prompt entry, (c) Mask creation, and (d) Text manipulation.

Table 1: Quality evaluation of AI-manipulated invoice and receipt images.

Dataset	PSNR	SSIM	LPIPS
Invoice Images	41.5120	0.9956	0.0024
Receipt Images	35.4529	0.9938	0.0042

text is selected, the binary mask is generated that highlights the region to be manipulated. Finally, the index number of the target text, the prompt of the new text, and the generated mask are used to guide TextDiffuser-2 [9] generative AI model in manipulating the input invoice. The proposed dataset comprises 5,324 images along with their ground truths and manipulated images, where 4,977 are invoice-based images sourced from the high-quality invoice dataset [1] and 347 are receipt-based images [27].

To evaluate the quality of AI-manipulated invoices and receipts, we computed three metrics, including PSNR, SSIM, and LPIPS. These metrics are used to measure the general visual quality, structural, and perceptual similarity. As depicted in Table 1, the manipulated invoices and receipts achieved high PSNR and SSIM scores and very low LPIPS scores, which suggest the manipulations introduce high perceptual realism with no noticeable artifacts.

4.2 Detection and localisation

To detect AI-manipulated receipts and invoices, we adopted an existing state-of-the-art technique FUSION++ [2], which is mainly designed to detect and localise generative AI manipulations. As illustrated in Figure 2, it employs an encoder-decoder based architecture where the encoder network extracts features from the input receipt $\mathbf{I}$, while the decoder detects AI manipulations and generates a localisation map $\mathbf{M}$. The encoder processed the input through the RGB (Red, Green and Blue color channels) path and the filters path, incorporating Histogram of Oriented Gradients (HOG), Noisprint++, Bayar, and Steganalysis Rich Model (SRM) filters to capture various visual manipulation artefacts such as noise and inconsistency in color, texture, and edges.

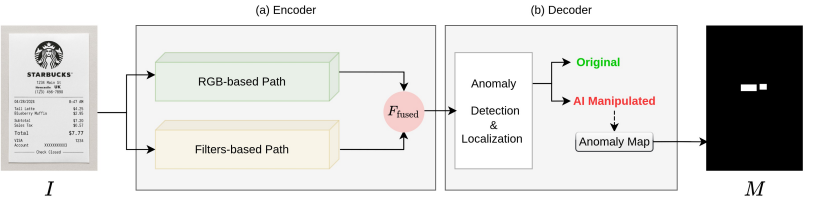

Figure 2: High-level block diagram of the FUSION++ based detection and localisation.

Table 2: Detection performance across the receipt and invoice datasets.

Method	**Receipt Dataset** (347 sample)		**Invoice Dataset** (4,977 sample)		**Average**	
	ACC	AUC	ACC	AUC	ACC	AUC
Proposed approach with FUSION++	**0.5288**	**0.5357**	**0.5008**	**0.5099**	**0.5148**	**0.5228**
HiFi-IFDL [15]	0.5000	0.4946	0.5000	0.5026	0.5000	0.4986

First, the encoder processed the inputs from the RGB and the filters paths separately through **Conv Blocks** to extract the initial feature maps. The maps from the filters (HOG, Noisprint++, Bayar, and SRM) are fused and refined through another **Conv Blocks**. Then, the extracted feature maps from both paths are enhanced through a shifted window-based multi-head self-attention **(SW-MHSA) Blocks** to focus on the most meaningful features and fused to form a final feature map $\mathbf{F_{fused}}$. The decoder processed the fused feature map to detect whether the receipt is original or manipulated and to localise the manipulated region.

As depicted in Table 2, we compare the detection performance of FUSION++ [7], adopted in this work and HiFi-IFDL [15], another recent non-finance generic media manipulation detector, using the Accuracy (ACC) and Area Under the Curve (AUC) metrics, with FUSION++ based approach surpassing HiFi-IFDL on both datasets. The best results are marked in bold. Figure 3 shows example predicted localisation maps from each category for manipulated invoices and receipts, respectively.

4.3 Discussion

Existing models such as Okamoto et al [13] focused on producing detection systems for traditional manipulations such as splicing, this system is the first known model that is capable of detecting manipulations produced using machine learning based techniques, which are typically harder to detect due to the ability of machine learning models to be able to match the style, font, layout and background noise in the images. In addition to this it can be determined that generic image manipulation detection frameworks and systems often struggle to detect smaller manipulations, which are likely to be commonplace on manipulated receipts. Therefore we can determine that this model provides a good benchmark for future models to build on and contribute to assisting in anti fraud efforts in the age of generative-AI.

5 Conclusion

The rise of AI-generated fake receipts and invoices signals a shift in financial document fraud. No longer limited by human effort or basic manipulation tools, fraudsters can now

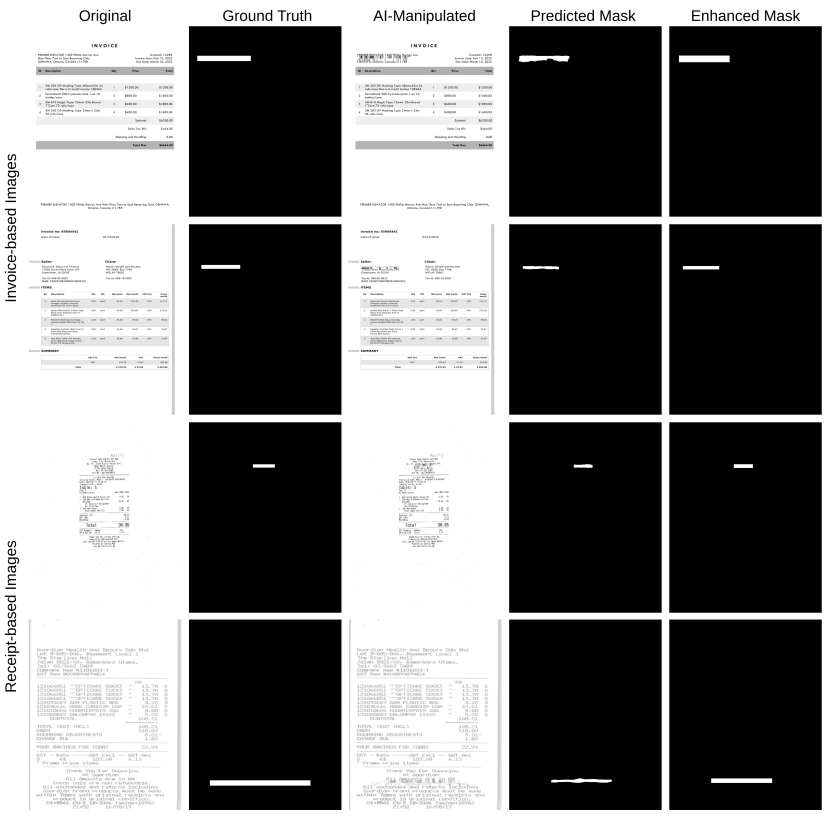

Figure 3: Prediction masks of AI-manipulated invoice and receipt using Fusion++.

create highly realistic, scalable forgeries that pose significant challenges to traditional detection methods. While some accounting professionals may view this as an evolution of an old problem, the advancement in automation, scale, and realism introduced by generative AI represents a fundamental threat to the integrity of financial systems. This paper has highlighted how these capabilities exploit weaknesses in current processes, underscoring the need for an updated response from the profession. The evidence reviewed makes clear that current safeguards are insufficient in isolation and require a proactive system-level design that embeds trust, traceability, and explainability at its core. The accounting industry stands at a crossroads, and its future resilience will depend on how successfully it can integrate emerging technologies such as AI and behavioural analytics into its core values of transparency and trust. The adoption of these innovations must be paired with strong education, ethical governance, and updated regulatory guidance.

In supporting a holistic framework, this paper also proposes techniques that allow for the creation of a pipeline of AI-generated invoices in addition to the identification of localised areas within an invoice in which AI-generated manipulation has occurred, which has been supported both by technical analysis and human-led review to prove its effectiveness. Future research will explore not only the technical solutions to AI-enabled fraud, but also how accountants deal with these risks in practice. Understanding behavioural, organisational, and cultural readiness will be essential for designing effective, usable safeguards.

References

[1] Osama Hosam Abdellatif and Abdelrahman Nader. High-quality invoice images for ocr, 2025. URL https://www.kaggle.com/dsv/11739593. Accessed: 3 September 2025.

[2] Lamyaa Aljuaid and Deepayan Bhowmik. FUSION++: A method to detect generative AI manipulated images. In *35th British Machine Vision Conference Workshop Proceedings, BMVC 2024 Workshops, Glasgow, UK, November 25-28, 2024*. BMVA Press, 2024. URL https://bmva-archive.org.uk/bmvc/2024/workshops/BMVAW/paper3.pdf.

[3] Author interviews. Interviews with accountants. Personal communication, 2025.

[4] Avey, C. How digital provenance preserves image integrity and security, 2025. URL https://www.secureworld.io/industry-news/digital-provenance-image-integrity-security. Accessed: 21 May 2025.

[5] Kar Balan, Robert Learney, and Tim Wood. A framework for cryptographic verifiability of end-to-end ai pipelines. In *Proceedings of the 10th ACM International Workshop on Security and Privacy Analytics*, IWSPA '25, page 49–59, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715013. URL https://doi.org/10.1145/3716815.3729011.

[6] Deepayan Bhowmik, Sabrina Caldwell, Jaime Delgado, Touradj Ebrahimi, Nikolaos Fotos, Xiaojun Gu, Ziyuan Hu, Xin Kang, Fernando Pereira, Leonard Rosenthol, et al. An international standard for assessing trustworthiness in media. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 3799–3805. IEEE, 2024.

[7] Bipartisan Policy Center. Ai use in tax administration, 2024. URL https://bipartisanpolicy.org/blog/tax-administration-and-ai/. Accessed: 29 May 2025.

[8] Booz Allen Hamilton. Deepfakes targeting benefits with ai-generated claims, 2025. URL https://www.boozallen.com/insights/ai-research/deepfakes-targeting-benefits-with-ai-generated-claims.html. Accessed: 18 July 2025.

[9] Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser-2: Unleashing the power of language models for text rendering. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part V*, page 386–402, Berlin, Heidelberg, 2024. Springer-Verlag. ISBN 978-3-031-72651-4. URL https://doi.org/10.1007/978-3-031-72652-1_23.

[10] European Commission. Vat in the digital age (vida), 2024. URL https://taxation-customs.ec.europa.eu/taxation/vat/vat-digital-age-vida_en. Accessed: 21 May 2025.

[11] FBI IC3. Criminals use generative artificial intelligence to facilitate financial fraud. Technical report, Federal Bureau of Investigation, 2024. URL https://www.ic3.gov/PSA/2024/PSA241203. Accessed: 29 May 2025.

[12] Financial Reporting Council. Frs 102: The financial reporting standard applicable in the uk and republic of ireland. Technical report, Financial Reporting Council, Sep 2024. URL https://media.frc.org.uk/documents/FRS_102_September_2024_tmKYWO6.pdf. Accessed: 8 September 2025.

[13] R. Gricius and I. Belovas. On the generation of synthetic invoices for training machine learning models. *IEEE Access*, 2025. URL https://doi.org/10.1109/ACCESS.2025.3555155.

[14] M Grosse. Can you spot a financial fake? how ai is raising risks of billing fraud. TechXplore, 10 April, 2025. URL https://techxplore.com/news/2025-04-financial-fake-ai-billing-fraud.html. Accessed: 12 July 2025.

[15] Xiao Guo, Xiaohong Liu, Zhiyuan Ren, Steven Grosz, Iacopo Masi, and Xiaoming Liu. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3165, 2023.

[16] S. Gupta and S. Kumar. Synthetic data generation using generative ai to combat identity fraud and enhance global financial cybersecurity frameworks. *International Journal of Scientific Research and Modern Technology*, 4(2):1–19, 2025. URL https://www.ijsrmt.com/index.php/ijsrmt/article/view/327.

[17] J. Kokina and T. H. Davenport. The emergence of artificial intelligence: Implications for data-driven auditing. *Journal of Emerging Technologies in Accounting*, 14(1):115–122, 2017. URL https://doi.org/10.5281/zenodo.14928919.

[18] M. Mulyadi and Y. Anwar. Business school teaching case study: taking accountancy from spreadsheets to ai. Financial Times, 2025. URL https://www.ft.com/content/bd9c415f-cab5-4ae1-8bf2-a17c57f9b5db. Accessed: 11 August 2025.

[19] Meharuniza Nazeem, Anitha R, Navaneeth S, and Rajeev R. R. Open-source OCR libraries: A comprehensive study for low resource language. In Sobha Lalitha Devi and Karunesh Arora, editors, *Proceedings of the 21st International Conference on Natural Language Processing (ICON)*, pages 416–421, AU-KBC Research Centre, Chennai, India, December 2024. NLP Association of India (NLPAI). URL https://aclanthology.org/2024.icon-1.48/.

[20] J. R. Nicolás-Carlock and I. Luna-Pla. Organized crime behavior of shell-company networks in procurement: Prevention insights for policy and reform. *Trends in Organized Crime*, 2023. URL https://doi.org/10.1007/s12117-023-09499-w.

[21] OECD. Technology tools to tackle tax evasion and tax fraud. Technical report, Organisation for Economic Co-operation and Development, Paris, 2017. URL https://doi.org/10.1787/g2g77afa-en. Accessed: 11 August 2025.

[22] Y Okamoto, G Osada, I Yahiro, R Hasegawa, P Zhu, and H Kataoka. Image generation and learning strategy for deep document forgery detection. *arXiv preprint arXiv:2311.03650*, 2023. URL https://arxiv.org/abs/2311.03650. Accessed: 7 June 2025.

[23] C. Rollet. Chatgpt's new image generator is really good at faking receipts. TechCrunch, 31 March, 2025. URL https://techcrunch.com/2025/03/31/chatgpts-new-image-generator-is-really-good-at-faking-receipts/. Accessed: 12 July 2025.

[24] M. Schreyer, T Sattarov, C Schulze, B Reimer, and D Borth. Detection of accounting anomalies in the latent space using adversarial autoencoder neural networks. *arXiv preprint arXiv:1908.00734*, 2019. URL https://doi.org/10.48550/arXiv.1908.00734. Accessed: 7 June 2025.

[25] U.S. Government Accountability Office. Fraud risk management: 2018–2022 data show federal government loses an estimated $233 billion to $521 billion annually to fraud, based on various risk environments. GAO Technical Report GAO-24-105833, U.S. Government Accountability Office (GAO), 2024. URL https://www.gao.gov/products/gao-24-105833.

[26] D. F. Wahid and E. Hassini. An augmented ai-based hybrid fraud detection framework for invoicing platforms. *Applied Intelligence*, 54:1297–1310, 2024. URL https://doi.org/10.1007/s10489-023-05223-x.

[27] Niansong Zhang, Songyi Yang, and Shengjie Xiu. Icdar 2019 robust reading challenge on scanned receipts ocr and information extraction. https://github.com/zzzDavid/ICDAR-2019-SROIE, 2019. Accessed: 3 September 2025.

[28] L. Zhao, C. Chen, and J. Huang. Deep learning-based forgery attack on document images. *arXiv preprint arXiv:2102.00653*, 2021. URL https://arxiv.org/pdf/2102.00653. Accessed: 7 June 2025.

Acknowledgment

This work is supported by funding from Sage Plc, through a PhD studentship of Mr Thornton and a collaborative project with Newcastle University. Ms Aljuaid and Mr Chapman are supported by PhD studentships via the Government of Saudi Arabia and Airbus Defence and Space, through the Engineering and Physical Science Research Council (EPSRC Reference: EP/S023577/1), within the Geospatial Systems CDT at Newcastle University, respectively.