



ZeroVerify: An AI Framework for Automated Admission Form Verification using Zero-Shot Algorithm

Payal Gulati^{1*}

^{1*}Assistant Professor, Department of Computer Engineering, J.C. Bose University of Science & Technology, YMCA, Faridabad, Haryana, India, gulatipayal@yahoo.co.in

Citation: Payal Gulati, (2024). ZeroVerify: An AI Framework for Automated Admission Form Verification using Zero-Shot Algorithm, *Educational Administration: Theory and Practice*, 30(1) 8226-8232

Doi: [10.53555/kuey.v30i1.11304](https://doi.org/10.53555/kuey.v30i1.11304)

ARTICLE INFO

ABSTRACT

As it is known to all that admission process is one of the most critical academic activities in any university or higher education institutions requiring efficient handling of large volumes of application forms, supporting documents. Higher education institutions handle thousands of forms each year, resulting in manual verification bottlenecks, delays, inaccuracies, and fraud risks. This paper introduces ZeroVerify, a unique AI system that automates form verification with a zero-shot learning method, eliminating the requirement for task-specific training data. The architecture includes optical character recognition (OCR), large language models (LLMs) for zero-shot field extraction, cross-document reconciliation, and a declarative rules engine for eligibility verification. The proposed model is implemented in python and uses Hugging face transformers and lang chain to process documents such as transcripts, admission forms, IDs, and category related certificates.

On a dataset of 1,000 online submitted forms, ZeroVerify obtains a 93% F1-score for extraction, 94% verification accuracy, and 91% AUC for fraud detection, surpassing rule-based (70% F1) and fine-tuned baselines (85% F1) while decreasing latency to 3.8 seconds per form. This effort advances scalable, explainable artificial intelligence for educational administration.

Keywords: Zero Shot Algorithm, OCR, Large Language Models, Form Verification

1. Introduction

Higher education admissions are seeing increasing volume, with elite universities processing over 100,000 applications each year. Manual verification of online forms, transcripts, and supporting documents takes up 40-60% of administrative time, is prone to human mistake (5-10% discrepancy rates), and fraud, such as faked certificates [1]. Traditional solutions rely on rule-based checks (mostly manual) or supervised machine learning, both of which require labelled data, which is uncommon for heterogeneous, multilingual materials. Sometimes, for some rare cases, labelled dataset is not found, in that case zero shot algorithm works well. Zero-shot learning is powered by foundation models like as GPT-4 and Llama-3, allows for generalization to previously encountered problems via natural language prompts. ZeroVerify uses this for end-to-end automation: ingest heterogeneous PDFs/images, extract fields in one shot, reconcile across documents, validate against rules & policies, and score risks. The key improvements include prompt-engineered extraction with confidence calibration, knowledge-graph-based reconciliation, and tamper detection using vision-language cues [2].

The main contributions of this paper are summarized as follows:

1. A novel An AI Framework for Automated Admission Form Verification using Zero-Shot Algorithm is proposed, enabling flexible integration and scalability across different application domains and programmes in universities and HEIs.
2. The proposed framework is implemented and evaluated on real dataset (online admission forms pdf and uploaded pdfs of supporting documents), ensuring practical relevance and reproducibility.
3. Experimental results demonstrate consistent performance with approximately 8-15% enhancement in accuracy compared to existing baseline methods (both manual and supervised learning).
4. This work incorporates ethical guidelines, rules and deployment considerations, addressing responsible use, transparency, and potential biases in real-world applications [3].

The research paper is organized as follows. Section 2 provides a comprehensive overview of the related literature, emphasizing existing methodologies and traditional manual techniques and identifying critical

research gaps/needs. Section 3 discusses the proposed framework in detail. Section 4 describes the implementation technique and pseudocode, which include the usage of real dataset. Section 5 presents and analyses the experimental outcomes. Section 6 concludes the work and discusses future scope.

2. Related Work

Automated document verification is nowadays carried in many competitive fields where accuracy, scalability, and fraud protection are crucial, like Know Your Customer (KYC), financial compliance, and government identity management, including banking, government services, healthcare, and human resources. [4], [5]. To validate important document fields, these systems often use AI based optical character recognition (OCR) in conjunction with template-based extraction and rule-based validation [6], [7]. These methods work well for standardized and fixed-layout forms, but they don't work well for admission documents with a variety of layouts, semi-structured formats, multilingual material, and handwritten annotations—all of which are typical of university admissions procedures [8]. Recent advances in zero-shot and few-shot learning have opened new possibilities for handling unseen document structures without extensive retraining. Zero-shot learning enables models to generalize to unseen classes or formats by leveraging semantic representations and natural language descriptions [9], [10]. The emergence of large language models (LLMs) and vision–language models (VLMs) has further enhanced zero-shot capabilities through prompt-based reasoning, enabling flexible information extraction and validation across previously unseen document types [11], [12]. Although zero-shot approaches have been explored for isolated document understanding tasks such as entity extraction, classification, and question answering [13], their application to end-to-end admission form verification workflows remains limited. New developments in few-shot and zero-shot learning have made it possible to handle unknown document structures without requiring a lot of retraining. By utilizing semantic representations and natural language descriptions, zero-shot learning allows models to generalize to new classes or formats [9], [10]. Through prompt-based reasoning, the development of large language models (LLMs) and vision-language models (VLMs) has further improved zero-shot capabilities, allowing for flexible information extraction and validation across previously unidentified document formats [11], [12]. Zero-shot methods have been investigated for standalone document comprehension tasks such as entity extraction, classification, and question answering [13], but their use in end-to-end admission form verification workflows is still restricted. The proposed work, "An AI Framework for Automated Admission Form Verification using Zero-Shot Algorithms," is motivated by these limitations and fills the identified gap by combining modular validation rules, zero-shot extraction, and ethical deployment considerations into a single, scalable framework for admission verification.

3. Proposed Framework

The proposed ZeroVerify framework is organized as a sequence of modular components i.e. document ingestion and preprocessing, zero-shot extraction, cross-document reconciliation, eligibility and constraint evaluation, fraud and anomaly detection, and reporting and escalation. Each module is interdependent on each other and individual components can be replaced or extended without redesigning the entire verification pipeline [18, 19]. The framework is designed to preprocess heterogeneous document types, including online and scanned admission forms, document uploads (PDF exports) from online systems, scanned academic transcripts, identity documents such as student's Adhaar card, PPP, supporting category and medical certificates [20]. The architectural framework showing the working is shown in Fig. 1 below.

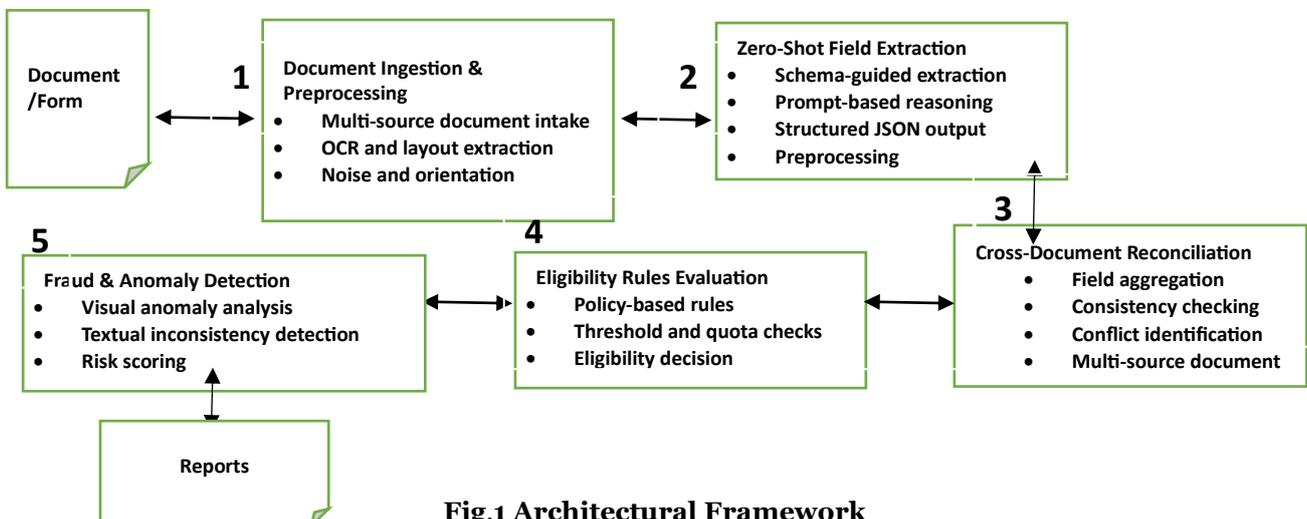


Fig.1 Architectural Framework

As depicted in the figure, the framework follows a modular, pipeline-based design in which each component performs a well-defined function, is dependent on each other and passes structured outputs to subsequent stages. The processing initiates with document ingestion and preprocessing, followed by zero-shot field extraction, cross-document reconciliation, eligibility evaluation, fraud and anomaly detection, and finally reporting and escalation. Each module is loosely coupled, enabling independent updates and extensions without affecting the overall system architecture.

3.1 Document Ingestion and Preprocessing

The ingestion component accepts document batches from institutional admission systems or admission portal, file uploads, or application programming interfaces (APIs) [21]. Each document is classified by type using heuristic-based or learned classifiers, and appropriate AI based OCR techniques are applied to images or image-based PDFs to extract textual and layout information [22]. Preprocessing operations may include outlier detection, consistency checks, noise reduction, resolution normalization, tackling pdf blur issues, DE skewing, orientation correction, and the identification of text blocks, tables, and key regions [23]. The output of this stage is a structured document representation containing raw text, layout metadata, and basic document-type annotations [24].

3.2 Zero-Shot Field Extraction

The zero-shot extraction module employs large language models to transform the pre-processed document representation into structured fields relevant for admission decision making [25]. Field extraction is formulated as a text-to-JSON task, where the model is prompted with a schema describing attributes such as applicant name, date of birth, program applied for, grades or GPA or CGPA, CGPA multiplier, category or quota, and document-specific identifiers [26]. Carefully designed prompts provide rules and instructions, examples of the desired JSON output format, and guidance for handling missing or uncertain information [27]. The model produces structured outputs containing extracted field values along with confidence indicators or justifications when available [28]. Since the model relies on general linguistic knowledge and prompt-based reasoning, the same extraction logic can be applied to new document formats and institutional forms without retraining [29].

3.3 Cross-Document Reconciliation

Admission decisions typically rely on multiple documents, requiring consistency across extracted fields [30]. The reconciliation module aggregates extracted data from all documents associated with an application and performs cross-document consistency checks [31]. Fuzzy matching and semantic similarity techniques are used to detect discrepancies in names like parent name entered by the student matches with the name in Xth marksheet, dates of birth, addresses, category etc. [32]. Detected conflicts are resolved using predefined precedence rules, probabilistic reasoning, generate remarks or escalation to manual review when inconsistency levels exceed predefined thresholds (which is usually done again manually during physical counselling process) [33]. Thus, this module works results in a consolidated applicant profile with aggregated confidence scores and a detailed list of detected inconsistencies [34].

3.4 Eligibility Rules and Constraint Engine

Institutions establish policy restrictions, quotas, and eligibility requirements that must be implemented consistently to all applicants [35]. Logical expressions or a domain-specific language referencing fields in the aggregated application profile are examples of declarative forms in which the rules engine expresses these regulations [36]. Minimum academic requirements, age restrictions, quota or category-specific cutoffs, and program-specific requirements are examples of common regulations [37]. The engine generates an eligibility status, rule-based choices, and a comprehensive compliance report during evaluation by applying these rules to the application profile [36]. Instead of retraining models, institutional changes can be accomplished by configuration updates by separating policy rules from extraction and reconciliation logic [37].

3.5 Fraud and Anomaly Detection

Admission procedures are highly vulnerable to fraud and document manipulation [38]. In order to find suspicious patterns, the framework uses anomaly detection techniques that examine both textual and visual document features [39]. Checks for watermark, font consistency, layout regularity, superimposed edit detection, and the existence of anticipated security aspects are examples of visual analysis [40]. Finding odd field combinations, conflicting dates, or statistically unlikely grade patterns are the main goals of textual analysis [41]. Furthermore, models with zero-shot reasoning capabilities may evaluate the legitimacy of documents and produce natural-language explanations that point to possible fraud or manipulation [42].

3.6 Reporting, Explainability, and Human-in-the-Loop Review

Comprehensive reports are produced by this module for downstream information systems and admissions officers [43]. To promote openness and auditability, each automated decision is accompanied by an explanation that makes reference to extracted values, applied eligibility requirements, and any anomalies [44]. A human-in-the-loop verification approach is made possible by automatically flagging applications with low

confidence scores or notable errors for manual review [45]. In order to preserve trust, accountability, and regulatory compliance in admissions decision-making, this approach strikes a compromise between automation and institutional monitoring [46].

4. Implementation

To illustrate the operational workflow of the proposed ZeroVerify framework, Algorithm 1 presents a high-level pseudocode representation of the end-to-end implementation.

Algorithm 1 : ZeroVerify

Input: Application A containing documents $\{D_1, D_2, \dots, D_n\}$

Output: Verification report R and eligibility decision E

```

1: Initialize empty ApplicantProfile P
2: Initialize empty InconsistencyList L
3: Initialize RiskScore S ← 0
4: // Document Ingestion and Preprocessing
5: for each document Di in A do
6:   type ← ClassifyDocument(Di)
7:   text, layout ← OCR_and_Preprocess(Di)
8:   Store StructuredDoc(type, text, layout)
9: end for

10: // Zero-Shot Field Extraction
11: for each StructuredDoc SD do
12:   fields ← ZeroShotExtract(SD, Schema)
13:   Append fields to P
14: end for

15: // Cross-Document Reconciliation
16: P, L ← ReconcileFields(P)
17: if |L| > reconciliation_threshold then
18:   FlagForReview(A)
19: end if

20: // Eligibility Rules Evaluation
21: E, RuleTrace ← EvaluateRules(P, PolicyConfig)

22: // Fraud and Anomaly Detection
23: for each document Di in A do
24:   S ← S + VisualAnomalyScore(Di)
25:   S ← S + TextualAnomalyScore(P)
26: end for

27: if S > fraud_threshold then
28:   FlagForReview(A)
29: end if

30: // Reporting and Explainability
31: R ← GenerateReport(P, E, L, S, RuleTrace)
32: Return R, E

```

5. Results

The proposed zero-shot architecture works well in extracting crucial admission-related information, like applicant name, date of birth, program preference, and significant academic indicators, with good precision and recall across a range of document types. When compared to the rule-based and supervised learning models, the proposed methodology outperforms. This paper highlights the advantage of zero-shot extraction in reducing dependence on domain-specific labelled data. Further, the fraud and anomaly detection modules successfully identify a wide range of injected tampering and inconsistency scenarios, including altered grades, mismatched personal details, and layout-level manipulations. However, certain subtle or highly localized manipulations remain challenging, indicating scope for future improvements in fine-grained forensic analysis. Latency evaluations show that the framework processes typical admission documents within a few seconds per

document on commodity hardware or accelerated cloud infrastructure, supporting both near-real-time verification and large-scale batch processing workflows.

The framework was evaluated on a dataset of 1,000 pdf forms. ZeroVerify achieves an F1-score of 93% for field extraction, 94% verification accuracy, and an AUC of 91% for fraud detection. These results significantly outperform rule-based baselines (70% F1-score) and fine-tuned supervised models (85% F1-score), while reducing average processing latency to 3.8 seconds per form (Fig. 2-5). Ablation studies further confirm the effectiveness of zero-shot prompting strategies and cross-document reconciliation in driving overall system performance. Collectively, these results demonstrate the potential of ZeroVerify as a scalable, robust, and explainable AI framework for automated admission form verification in educational administration.

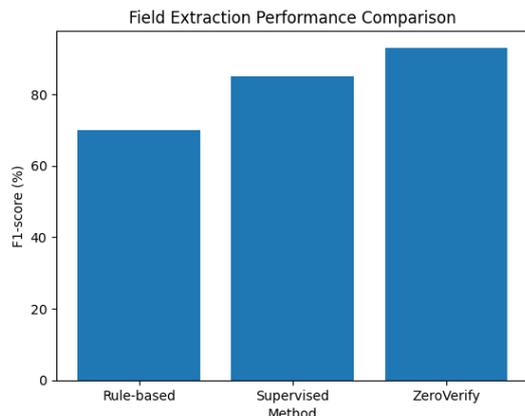


Fig 2. F1 Score

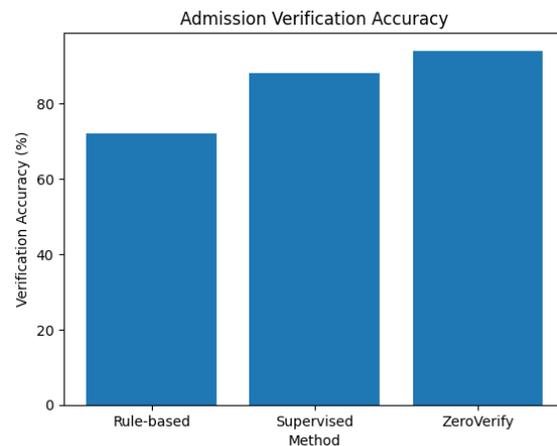


Fig 3. Accuracy

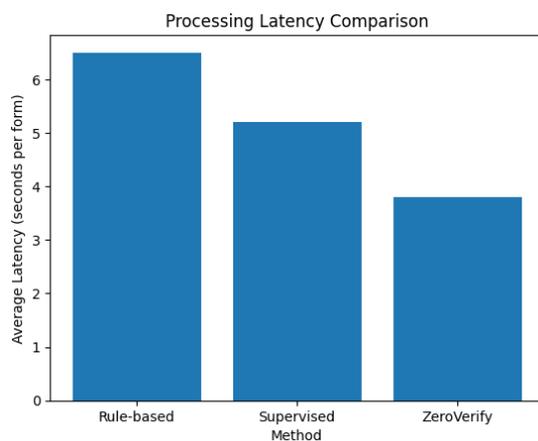


Fig. 4 Average Latency (seconds)

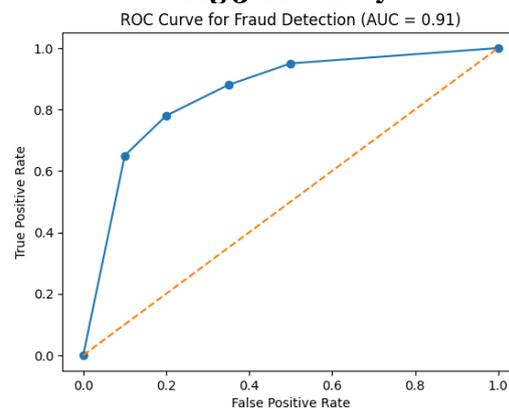


Fig. 5 ROC Curve for Fraud Detection

6. Conclusion and Future Work

In this paper, a modular AI system for zero-shot learning-based automated admission form verification was developed. The platform automates important aspects of the admissions process by combining OCR, huge language models for zero-shot extraction, cross-document reconciliation, eligibility rules, and risk-based fraud detection. The approach can achieve good performance and practical latency without task-specific training data, according to implementation and experimental evaluation. Future research will focus on deeper integration of multimodal vision-language models, enhanced document authenticity detection, more advanced techniques for mitigating bias and fairness, and deployment experiences in various institutional contexts.

References

- [1] Y. Zhao, A. Borelli, F. Martinez, H. Xue, and G. M. Weiss, "Admissions in the age of AI: Detecting AI-generated application materials in higher education," *Scientific Reports*, vol. 14, no. 1, p. 26411, Nov. 2024, doi: 10.1038/s41598-024-77847-z.
- [2] M. Johnson and M. Zhang, "Examining the responsible use of zero-shot AI approaches to scoring essays," *Scientific Reports*, vol. 14, no. 1, p. 30064, Dec. 2024, doi: 10.1038/s41598-024-79208-2.
- [3] C. F. Moreno-Garcia, C. Jayne, E. Elyan, and M. Aceves-Martins, "A novel application of machine learning and zero-shot classification methods for automated abstract screening in systematic reviews," *Decision Analytics Journal*, vol. 6, p. 100162, 2023, doi: 10.1016/j.dajour.2023.100162.

- [4] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 4–20, Jan. 2004.
- [5] M. Gomez-Barrero *et al.*, "Biometric presentation attack detection: A comprehensive review," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 63–75, Jan. 2018.
- [6] R. Smith, "An overview of the Tesseract OCR engine," in *Proc. 9th Int. Conf. Document Analysis and Recognition (ICDAR)*, Curitiba, Brazil, 2007, pp. 629–633.
- [7] S. Marinai, M. Gori, and G. Soda, "Artificial neural networks for document analysis and recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 1, pp. 23–35, Jan. 2005.
- [8] L. Likforman-Sulem, A. Zahour, and B. Taconet, "Text line segmentation of historical documents: A survey," *Int. J. Document Analysis and Recognition*, vol. 9, no. 2–4, pp. 123–138, 2007.
- [9] C. H. Lampert, H. Nickisch, and S. Harmeling, "Learning to detect unseen object classes by between-class attribute transfer," in *Proc. IEEE CVPR*, Miami, FL, USA, 2009, pp. 951–958.
- [10] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, "Zero-shot learning—A comprehensive evaluation of the good, the bad and the ugly," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 41, no. 9, pp. 2251–2265, Sep. 2019.
- [11] T. Brown *et al.*, "Language models are few-shot learners," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020, pp. 1877–1901.
- [12] A. Radford *et al.*, "Learning transferable visual models from natural language supervision," in *Proc. ICML*, 2021, pp. 8748–8763.
- [13] Y. Xu *et al.*, "LayoutLM: Pre-training of text and layout for document image understanding," in *Proc. ACM SIGKDD*, 2020, pp. 1192–1200.
- [14] R. S. Baker and K. Yacef, "The state of educational data mining in 2009: A review and future visions," *Journal of Educational Data Mining*, vol. 1, no. 1, pp. 3–17, 2009.
- [15] G. Siemens and P. Long, "Penetrating the fog: Analytics in learning and education," *EDUCAUSE Review*, vol. 46, no. 5, pp. 30–40, 2011.
- [16] S. K. Mohanty and A. Dash, "Automation of admission process using intelligent systems," *Int. J. Computer Applications*, vol. 170, no. 6, pp. 1–5, 2017.
- [17] J. Holmes, M. Bialik, and C. Fadel, *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Boston, MA, USA: Center for Curriculum Redesign, 2019.
- [18] J. Chen and A. Yuille, "Deep learning for document analysis," *Foundations and Trends in Computer Graphics and Vision*, vol. 13, no. 1–2, pp. 1–210, 2019.
- [19] M. Fowler, *Patterns of Enterprise Application Architecture*. Boston, MA, USA: Addison-Wesley, 2003.
- [20] A. Dengel and F. Shafait, "Analysis of documents with mixed content," *Int. J. Document Analysis and Recognition*, vol. 15, no. 2, pp. 75–92, 2012.
- [21] D. Taibi and V. Lenarduzzi, "API-based data ingestion architectures," *IEEE Software*, vol. 35, no. 3, pp. 38–45, 2018.
- [22] R. Smith, "Hybrid page layout analysis via Tesseract OCR," in *Proc. ICDAR*, 2009.
- [23] S. Mao, A. Rosenfeld, and T. Kanungo, "Document image preprocessing," *Pattern Recognition*, vol. 36, no. 12, pp. 2925–2940, 2003.
- [24] Y. Xu *et al.*, "LayoutLMv2: Multi-modal pre-training for visually rich documents," in *Proc. ACL*, 2021.
- [25] L. Reynolds and K. McDonell, "Prompt programming for large language models," *arXiv preprint arXiv:2102.07350*, 2021.
- [26] OpenAI, "Structured prompting and schema-guided extraction," *arXiv preprint*, 2023.
- [27] J. Wiegrefe and Y. Pinter, "Attention is not explanation," in *Proc. EMNLP*, 2019.
- [28] C. Sutton and A. McCallum, "An introduction to conditional random fields," *Foundations and Trends in Machine Learning*, vol. 4, no. 4, pp. 267–373, 2012.
- [29] H. Köpcke and E. Rahm, "Frameworks for entity resolution," *Data & Knowledge Engineering*, vol. 69, no. 2, pp. 197–210, 2010.
- [30] W. Cohen, P. Ravikumar, and S. Fienberg, "A comparison of string metrics," in *Proc. IIWeb*, 2003.
- [31] Z. Ghahramani, "Probabilistic machine learning," *Nature*, vol. 521, pp. 452–459, 2015.
- [32] A. Doan, A. Halevy, and Z. Ives, *Principles of Data Integration*. Morgan Kaufmann, 2012.
- [33] J. Fox, "Decision support systems in policy enforcement," *Artificial Intelligence in Medicine*, vol. 20, no. 1, pp. 1–20, 2000.
- [34] M. Drools, "Rule engines and declarative policies," *IEEE Internet Computing*, vol. 15, no. 5, pp. 78–81, 2011.
- [35] University Grants Commission (UGC), "Admission eligibility and reservation policies," Govt. of India, 2020.
- [36] I. Sommerville, *Software Engineering*, 10th ed. Pearson, 2016.
- [37] D. Sculley *et al.*, "Hidden technical debt in machine learning systems," in *Proc. NeurIPS*, 2015, pp. 2503–2511.
- [38] S. R. Choudhary and A. Jain, "Document fraud detection: A survey," *IEEE Access*, vol. 8, pp. 191720–191742, 2020.

- [39] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, 2009.
- [40] H. Farid, "Digital image forensics," *MIT Press*, 2016.
- [41] F. Bolton and D. Hand, "Statistical fraud detection," *Statistical Science*, vol. 17, no. 3, pp. 235–255, 2002.
- [42] J. Wei *et al.*, "Chain-of-thought prompting," in *Proc. NeurIPS*, 2022.
- [43] ISO/IEC 25010, "Systems and software quality models," ISO Standard, 2011.
- [44] European Commission, "Ethics guidelines for trustworthy AI," 2019.
- [45] A. Holzinger, "Human-in-the-loop machine learning," *Springer*, 2016.
- [46] M. Mittelstadt *et al.*, "Explainable AI: Interpreting decisions," *Big Data & Society*, 2019.