JOURNAL OF GLOBAL ENGINEERING REVIEW ISSN: XXXX-XXXX

Semi-Supervised Learning Approach for Automated Sensitive Data Classification in Unstructured Text Documents

Juan Li[1] , Wenkun Ren[1.2] Xiaolan Wu[2]

[1] Shanghai Jiao Tong University Master of Science in Communication and Information Systems
[1.2]Information Technology and Management, Illinois Institute of Technology, Chicago, IL
[2] Northeastern University Computer Science

Abstract

Effective data protection demands accurate identification and categorization of sensitive information across organizational repositories. Manual classification methodologies introduce substantial temporal overhead while generating inconsistent taxonomic assignments that undermine governance frameworks. This research develops a semi-supervised learning architecture for automating sensitive data classification within unstructured textual environments, addressing the fundamental challenge of limited labeled training instances. We construct a probabilistic framework integrating self-training mechanisms with confidence-weighted label propagation, achieving 87.3% classification accuracy using merely 12% labeled data. The methodology applies natural language processing techniques for feature extraction from heterogeneous document formats including emails, reports, and collaborative workspace content. Our experimental evaluation across three organizational datasets demonstrates precision improvements of 14.6% over baseline supervised approaches under constrained labeling scenarios. We establish a four-tier sensitivity taxonomy aligned with regulatory compliance requirements and implement adaptive decision boundaries for human validation workflows. Performance analysis reveals consistent recall rates exceeding 82% across personally identifiable information, financial records, and proprietary intelligence categories. The framework reduces manual annotation requirements by 73% while maintaining classification fidelity sufficient for regulatory compliance auditing.

Keywords: *Sensitive Data Classification, Semi-Supervised Learning, Unstructured Text Processing, Privacy Compliance*

1. Introduction

1.1 Research Background and Motivation

Contemporary organizations generate unprecedented volumes of unstructured textual data through digital communication channels, collaborative platforms, and document management systems. This information landscape contains heterogeneous sensitive data elements demanding systematic identification and protection mechanisms. Extracting structured knowledge from unstructured text repositories presents fundamental challenges in information management[1]. The escalating regulatory environment mandates precise classification capabilities distinguishing public, internal, confidential, and restricted information categories. Traditional manual classification workflows impose substantial operational burdens while introducing human variability undermining consistency across enterprise data holdings.

Organizations confront dual imperatives of maintaining comprehensive data visibility while implementing granular access controls preventing unauthorized exposure. Personally identifiable information requires specialized handling protocols across storage, transmission, and processing operations[2]. The intersection of machine learning capabilities with natural language understanding creates opportunities for developing scalable classification architectures adapting to organizational contexts. Data sensitivity exists along continuous spectrums rather than discrete categorical boundaries, complicating automated classification efforts when training data availability remains constrained by annotation costs and domain expertise requirements[3].

Current commercial solutions predominantly employ rule-based pattern matching or supervised learning models demanding extensive labeled datasets. These methodologies struggle with vocabulary evolution, contextual ambiguity, and cross-domain generalization challenges. The financial and temporal investments required for comprehensive data annotation create adoption barriers, particularly for organizations lacking dedicated data governance infrastructure. Semi-supervised learning paradigms offer alternatives by exploiting unlabeled data abundance to augment limited labeled examples, potentially reducing annotation requirements while preserving classification performance.

1.2 Problem Statement and Challenges

Automated sensitive data classification confronts multiple technical and operational obstacles hindering practical deployment. The fundamental problem resides in developing classification models achieving acceptable accuracy with minimal labeled training data while generalizing across document types, authoring styles, and domain-specific terminology. Data privacy protection demands sophisticated definitional frameworks and technical implementations addressing diverse information categories[4]. Unstructured text presents inherent difficulties through variable formatting, mixed content types, embedded objects, and inconsistent structural organization. Sensitive information may appear in disparate contexts ranging from explicit identifiers to implicit references requiring semantic understanding.

The labeling bottleneck represents a critical constraint where subject matter experts must manually annotate documents according to sensitivity taxonomies. This process consumes substantial temporal resources while introducing subjective interpretation variations across annotators. Semi-supervised learning methods address labeled data scarcity through various algorithmic strategies[5]. Organizations possess extensive unlabeled document collections remaining unexploited in supervised learning frameworks. The challenge involves designing algorithms effectively utilizing unlabeled data to improve classification boundaries without compromising precision.

Regulatory compliance frameworks impose additional requirements for explainability and auditability, necessitating transparent decision-making processes enabling human oversight. Privacy regulations establish specific requirements for identifying and protecting particular data categories. General data protection regulation compliance demands systematic approaches to personally identifiable information management across distributed systems[6]. Classification systems must accommodate evolving threat landscapes, emerging data types, and dynamic organizational structures while maintaining operational efficiency. Deployment considerations include integration with existing security infrastructure, real-time processing capabilities, and minimal performance overhead on production systems.

1.3 Research Objectives and Contributions

This research develops a semi-supervised learning framework for automated sensitive data classification in unstructured text documents, addressing labeled data scarcity while maintaining classification fidelity. Semi-supervised learning combines labeled and unlabeled data to improve model generalization capabilities[7]. Our methodology integrates self-training algorithms with confidence-weighted label propagation, enabling effective learning from limited labeled instances. The framework applies natural language processing techniques for feature extraction from heterogeneous document formats, capturing semantic and syntactic patterns indicative of sensitive information categories.

We establish a four-tier sensitivity taxonomy encompassing public, internal, confidential, and restricted classifications aligned with organizational governance requirements and regulatory mandates. Generating linked data from unstructured text requires sophisticated extraction and transformation mechanisms[8]. The research implements adaptive decision boundaries enabling human-in-the-loop validation for borderline classifications, balancing automation efficiency with accuracy assurance. Our experimental evaluation across three organizational datasets demonstrates the framework's effectiveness under realistic operational constraints.

The primary contributions include: a semi-supervised learning architecture reducing labeled data requirements by 73% while achieving 87.3% classification accuracy; a probabilistic confidence estimation mechanism enabling selective human review of uncertain classifications; an evaluation protocol measuring performance across diverse document types and sensitivity categories; performance analysis revealing 14.6% precision improvements over supervised baselines with equivalent labeled data quantities. The framework provides organizations with scalable solutions for implementing data protection measures efficiently while maintaining classification quality sufficient for regulatory compliance auditing and governance reporting requirements.

2. Related Work and Background

2.1 Sensitive Data Classification Techniques and Applications

Sensitive data classification research encompasses diverse methodologies ranging from rule-based pattern matching to deep learning architectures. Data sensitivity exists along multidimensional spectrums incorporating legal, ethical, commercial, and technical considerations[9]. Early approaches employed regular expression patterns and keyword dictionaries for identifying specific information types like social security numbers, credit card numbers, and email addresses. These techniques achieve high precision for well-defined patterns but struggle with contextual variations and semantic ambiguity.

Machine learning-based classification methods apply supervised learning algorithms trained on labeled document collections. Privacy and security frameworks establish foundational definitions and technical approaches for protecting sensitive information[10]. Feature engineering strategies include bag-of-words

representations, term frequency-inverse document frequency weighting, and n-gram models capturing local contextual patterns. Support vector machines, random forests, and neural network architectures demonstrate effectiveness across various text classification tasks. Deep learning models employing convolutional and recurrent neural networks extract hierarchical feature representations from raw text, eliminating manual feature engineering requirements. Scanning mechanisms for personally identifiable information in electronic documents employ pattern recognition and contextual analysis techniques[11].

2.2 Semi-Supervised Learning Approaches in Text Classification

Semi-supervised learning addresses scenarios where labeled training data remains scarce while unlabeled data exists abundantly. Semi-supervised methodologies enable learning from both labeled and unlabeled instances through various algorithmic strategies[12]. Self-training approaches iteratively predict labels for unlabeled instances, adding high-confidence predictions to the training set. Co-training algorithms train multiple classifiers on different feature subsets, enabling each classifier to label instances for others. Graph-based methods construct similarity graphs over instances, propagating labels through graph structure based on edge weights reflecting instance similarity.

Generative models assume data generation processes, estimating parameters from both labeled and unlabeled data through expectation-maximization algorithms. Automated ontology construction from unstructured text documents employs semi-supervised techniques for knowledge extraction[13]. Transductive support vector machines extend maximum margin principles to unlabeled data, adjusting decision boundaries to maximize margins over entire datasets. Consistency regularization methods encourage prediction consistency under data perturbations, enforcing smooth decision boundaries. These approaches demonstrate effectiveness across text classification tasks, achieving performance comparable to supervised models while requiring substantially fewer labeled instances.

2.3 Privacy Regulations and Data Governance Requirements

Privacy regulations worldwide establish comprehensive requirements for sensitive data protection, imposing specific obligations on organizations collecting, processing, and storing personal information. Data security and privacy concepts encompass technical measures, organizational policies, and legal frameworks protecting information confidentiality and integrity[14]. The General Data Protection Regulation establishes stringent requirements for personal data processing within European Union jurisdictions, mandating lawful bases for processing, data subject rights, and breach notification obligations. The California Consumer Privacy Act grants consumers rights regarding personal information collection and usage, requiring businesses to implement technical and organizational measures ensuring data protection.

Regulatory frameworks necessitate systematic data classification enabling differential protection based on sensitivity levels. Mining applications for abnormal usage of sensitive data reveals privacy violations and security vulnerabilities requiring detection and remediation[15]. Organizations must implement data inventory and mapping processes identifying sensitive information across distributed systems. Classification systems must align with regulatory definitions while accommodating organizational risk tolerance and operational requirements[16]. Governance frameworks establish policies, procedures, and controls ensuring appropriate data handling throughout information lifecycles[17]. Compliance auditing demands verifiable classification processes with documented decision rationale and human oversight mechanisms[18]. Automated classification tools must provide sufficient transparency and accuracy for regulatory acceptance while maintaining operational efficiency at organizational scale[19].

3. Methodology

3.1 Data Collection and Preprocessing Pipeline

The data collection pipeline aggregates unstructured text documents from multiple organizational sources including email archives, shared drive repositories, collaborative workspace platforms, and document management systems. We construct three distinct datasets representing different organizational contexts[20]: a financial services corporation with 47,832 documents, a healthcare provider network with 32,156 documents, and a technology company with 54,921 documents[21]. Each dataset contains authentic organizational communications and documents with inherent sensitivity variations[22].

The preprocessing pipeline implements multi-stage transformation converting raw documents into structured feature representations suitable for machine learning algorithms[23]. Document parsing extracts textual content from diverse file formats including PDF, DOCX, TXT, HTML, and email formats (EML, MSG). Format-specific parsers handle embedded metadata, structural elements, and encoding variations. Text normalization applies lowercase conversion, Unicode standardization, and whitespace regularization while preserving structural boundaries between paragraphs and sections.

Tokenization segments continuous text into discrete linguistic units using whitespace and punctuation delimiters, with specialized handling for hyphenated terms, contractions, and domain-specific compound

expressions[24]. Stop word filtering removes high-frequency function words contributing minimal discriminative information for classification tasks[25]. The stop word list encompasses 421 terms including articles, prepositions, conjunctions, and common verbs. We retain domain-specific terminology and named entities despite high frequency values due to classification relevance[26].

Stemming algorithms reduce morphological variations to root forms, applying Porter stemming for English text with exception dictionaries preserving medical terminology, legal terms, and technical jargon where morphological normalization alters semantic meaning[27]. Named entity recognition identifies person names, organizations, locations, dates, monetary values, and other structured information types using conditional random field models trained on annotated corpora. Entity recognition achieves F1 scores of 0.89 for person names, 0.82 for organizations, and 0.91 for monetary values across evaluation datasets[28].

Feature extraction transforms preprocessed text into numerical representations for classification algorithms. We employ term frequency-inverse document frequency vectorization with vocabulary size limited to 15,000 terms ranked by document frequency. The TF-IDF weighting scheme applies:

$$\text{TF-IDF}(t,d) = \text{TF}(t,d) \times \log\left(\frac{N}{\text{DF}(t)}\right)$$

where TF(t,d) represents term frequency in document d, N denotes total document count, and DF(t) indicates document frequency for term t. Bigram features capture local context patterns, with bigram vocabulary limited to 8,000 most frequent pairs based on pointwise mutual information scores[29]. Document length normalization applies L2 normalization ensuring unit vector magnitudes preventing bias toward longer documents[30].

Manual labeling efforts engage domain experts for sensitivity assessment across four categories: Public (openly shareable), Internal (employee-accessible), Confidential (role-restricted), and Restricted (highly sensitive requiring specialized authorization). We establish labeling guidelines defining category characteristics with concrete examples and decision criteria. Inter-annotator agreement measured through Cohen's kappa achieves 0.78 across categories, indicating substantial agreement despite inherent classification ambiguity. The labeled dataset comprises 12% of total documents, distributed as: Financial (5,740 labeled), Healthcare (3,859 labeled), Technology (6,591 labeled).

Table 1: Dataset Characteristics and Label Distribution

Dataset	Total Documents	Labeled Documents	Public	Internal	Confidential	Restricted
Financial Services	47,832	5,740 (12.0%)	1,148	2,009	1,721	862
Healthcare Provider	32,156	3,859 (12.0%)	772	1,351	1,158	578
Technology Company	54,921	6,591 (12.0%)	1,318	2,307	1,977	989
Total	134,909	16,190 (12.0%)	3,238	5,667	4,856	2,429

3.2 Semi-Supervised Learning Framework for Sensitive Data Classification

The semi-supervised learning framework implements self-training with confidence-weighted label propagation, addressing limited labeled data availability while maintaining classification accuracy[31]. The architecture integrates base classifiers, confidence estimation mechanisms, and iterative refinement processes enabling progressive learning from unlabeled instances[32].

Base classifier selection evaluates multiple algorithms including logistic regression, random forests, and gradient boosting machines across labeled validation sets. Logistic regression with L2 regularization achieves baseline performance with interpretable coefficients indicating feature importance. The regularization parameter $\lambda = 0.01$ balances model complexity and training data fit, selected through cross-validation across $\lambda \in \{0.001, 0.01, 0.1, 1.0\}$. Random forests construct ensemble predictions from 200 decision trees with maximum depth 15, minimum samples per leaf 5, and bootstrap aggregation. Gradient boosting machines implement iterative tree construction with learning rate 0.1, maximum depth 8, and 150 boosting rounds. Validation set performance comparison selects gradient boosting machines as base classifiers, achieving 84.2% accuracy on held-out labeled data.

The self-training algorithm initializes with labeled dataset L and unlabeled dataset U, iteratively selecting high-confidence predictions from U for inclusion in expanded training set L'[33]. The confidence threshold τ

determines prediction acceptance criteria, balancing exploitation of model certainty and exploration of diverse instances[34]. We employ probability-based confidence estimation where classifier outputs probabilistic predictions P(y|x) for each class y given instance x. Confidence score computation applies:

$$\text{Confidence}(x) = \max\big(P(y|x)\big) - \text{Entropy}\big(P(y|x)\big)$$

where Entropy measures prediction uncertainty across classes. High maximum probability combined with low entropy indicates confident predictions[35]. The confidence threshold $\tau = 0.75$ admits predictions exceeding this combined metric, selected through validation set analysis of precision-recall tradeoffs at varying thresholds[36].

Label propagation extends self-training through graph-based semi-supervised learning, constructing similarity graphs connecting instances based on feature space proximity. Graph construction employs k-nearest neighbors with $k = 15$, computing edge weights through Gaussian kernel similarity:

$$w(i,j) = \exp\left(-\frac{|x_i - x_j|^2}{2\sigma^2}\right)$$

where $\sigma = 0.5$ controls neighborhood size. Label propagation iteratively updates unlabeled instance labels based on neighbor label distributions, implementing the update rule:

$$y_i^{(t+1)} = \alpha \cdot \frac{\sum_j \left(w(i,j) \cdot y_j^{(t)}\right)}{\sum_j w(i,j)} + (1-\alpha) \cdot y_i^{(0)}$$

where $\alpha = 0.8$ balances neighbor influence and initial predictions. Iteration continues until convergence with label changes below threshold $\varepsilon = 0.01$ or maximum iterations $T = 50$.

The integrated framework alternates self-training and label propagation phases, progressively expanding labeled set and refining classification boundaries. Algorithm 1 outlines the complete procedure:

Algorithm 1: Semi-Supervised Classification Framework

1. Initialize labeled set L, unlabeled set U

2. Train base classifier C on L

3. While |U| > 0 and iterations < max_iterations:

a. Predict labels for instances in U using C

b. Compute confidence scores for predictions

c. Select instances with confidence > τ

d. Add selected instances to L, remove from U

e. Construct similarity graph over L ∪ U

f. Apply label propagation on graph

g. Update labels for instances in U

h. Retrain classifier C on expanded L

4. Return final classifier C

The framework processes unlabeled instances in batches of 500 documents per iteration, balancing computational efficiency and label propagation effectiveness[37]. Each iteration requires classifier retraining, confidence assessment, graph construction, and label propagation, with average iteration time of 127 seconds on standard hardware (Intel Xeon E5-2680 v4, 64GB RAM). Total training time across all iterations averages 3.8 hours per dataset.

Table 2: Semi-Supervised Framework Configuration Parameters

Parameter	Value	Selection Rationale
Base Classifier	Gradient Boosting	Selected due to achieving the highest validation accuracy of

		84.2%, which demonstrates superior predictive performance compared to alternative classifiers evaluated.
Learning Rate	0.1	Determined through validation performance optimization, where this value balanced model convergence speed and the ability to avoid overshooting optimal weights during training.
Number of Trees	150	Established via convergence analysis on the validation set, indicating that model performance stabilized at this number of trees without significant improvement beyond this point, preventing unnecessary computational overhead.
Confidence Threshold (τ)	0.75	Selected based on precision-recall tradeoff analysis, striking a balance between minimizing false positives (high precision) and ensuring adequate coverage of true positives (reasonable recall) for the specific application requirements.
Graph k-NN (k)	15	Determined through neighborhood size sensitivity testing, where a k-value of 15 provided the optimal balance between capturing relevant local data structure and avoiding noise inclusion from overly distant neighbors.
Gaussian Kernel σ	0.5	Chosen based on distance distribution analysis of the data, ensuring that the kernel effectively weights nearby points while appropriately discounting the influence of more distant observations.
Label Propagation α	0.8	Selected via convergence stability assessment, with this value ensuring stable and reliable propagation of labels through the graph while maintaining a reasonable balance between trusting initial labels and propagated information.
Batch Size	500	Determined by computational efficiency consideration, balancing the need for frequent parameter updates (smaller batches) and leveraging hardware acceleration (larger

batches) to optimize training time and resource utilization.

3.3 Sensitivity Level Categorization and Labeling Strategy

Sensitivity level categorization establishes a four-tier taxonomy aligning with organizational governance requirements and regulatory compliance frameworks[38]. The taxonomy provides explicit definitions, concrete examples, and decision criteria for each category, ensuring consistent classification across diverse document types and organizational contexts[39].

Public category encompasses information intended for unrestricted external distribution with no confidentiality requirements[40]. Documents classified as Public include marketing materials, press releases, published research papers, public product documentation, and customer-facing website content[41]. Public information requires no access controls beyond standard publication channels, with disclosure presenting no organizational risk or competitive disadvantage[42].

Internal category covers information accessible to all organizational employees but restricted from external parties[43]. Internal documents include operational procedures, employee handbooks, internal newsletters, project status reports, and general business communications[44]. Internal classification requires authentication-based access controls preventing external access while permitting broad internal distribution[45]. Unauthorized external disclosure presents minimal risk beyond potential competitive intelligence concerns[46].

Confidential category encompasses sensitive business information requiring role-based access controls with disclosure limited to employees with legitimate business needs[47]. Confidential documents include financial statements, strategic plans, customer lists, proprietary methodologies, partner agreements, and detailed technical specifications[48]. Confidential classification implements fine-grained access controls based on organizational roles, departments, and project assignments[49]. Unauthorized disclosure risks competitive disadvantage, regulatory scrutiny, or customer relationship damage[50].

Restricted category represents highly sensitive information requiring specialized authorization with strictly limited access[51]. Restricted documents include personally identifiable information, protected health information, payment card data, trade secrets, merger and acquisition plans, and security vulnerability details. Restricted classification demands strongest access controls including multi-factor authentication, audit logging, encryption requirements, and data loss prevention monitoring[52]. Unauthorized disclosure creates substantial legal liability, regulatory penalties, or severe organizational harm[53].

The labeling strategy implements a hierarchical decision tree guiding human annotators through systematic assessment of document characteristics[54]. The decision tree evaluates multiple dimensions including: regulatory data type presence (PII, PHI, PCI), competitive sensitivity, legal privilege, confidentiality obligations, and potential disclosure impact[55]. Annotators answer structured questions at each decision node, traversing the tree until reaching terminal category assignment[56].

Active learning strategies optimize labeling efficiency by selecting most informative unlabeled instances for expert annotation[57]. Uncertainty sampling identifies instances with highest prediction entropy, indicating maximum model uncertainty. The entropy-based selection criterion:

$$\text{Uncertainty}(x) = -\Sigma\, P(y|x) \times \log(P(y|x))$$

selects instances maximizing this entropy measure for preferential labeling. Diversity sampling ensures selected instances represent feature space coverage, preventing redundant annotations for similar documents[58]. We implement clustering-based diversity selection, applying k-means clustering over unlabeled instances and selecting representatives from each cluster[59].

Query-by-committee approach trains multiple classifier variants on bootstrap samples of labeled data, selecting instances with maximum prediction disagreement among committee members[60]. The disagreement measure:

$$\text{Disagreement}(x) = (1/|C|) \times \Sigma(c \in C)\ 1[\text{argmax}\ P\_c(y|x) \neq \text{argmax}\ P\_\text{ensemble}(y|x)]$$

quantifies prediction variation across committee. High disagreement indicates informative instances clarifying classification boundaries.

Table 3: Sensitivity Category Definitions and Access Control Requirements

Category	Definition	Access Controls	Examples	Disclosure Risk

Public	Information that can be freely distributed to individuals or entities outside the organization without any limitations or restrictions.	No specific access controls other than the act of making the information publicly available (e.g., publishing on a website, distributing in public forums).	Marketing brochures, product catalogs, press releases, company newsletters intended for public consumption, and publicly available financial reports (if required by law).	Minimal, as the information is already intended for public knowledge and distribution, so unauthorized disclosure does not pose significant harm.
Internal	Information that is intended for access and use solely by employees of the organization and is not meant to be shared with individuals or entities outside the organization.	Basic authentication is required, typically in the form of a username and password, to ensure that only employees with valid organizational credentials can access the information.	Internal operational procedures, employee handbooks, internal memos, departmental meeting minutes, internal training materials, and internal announcements.	Low, as the information is primarily relevant to internal operations and its unauthorized disclosure to external parties may cause minor disruptions or embarrassment but not severe harm.
Confidential	Information that is sensitive and restricted to access by specific individuals within the organization based on their job roles, responsibilities, or specific need-to-know.	Fine-grained authorization mechanisms are implemented, which may include role-based access control (RBAC) where access permissions are assigned based on predefined roles, as well as additional access approval processes for certain highly sensitive confidential information.	Financial statements before public release, business strategies and plans, customer lists with sensitive details, proprietary research data, and confidential contracts.	Moderate to High, as unauthorized disclosure can lead to significant financial losses, damage to competitive advantage, harm to customer relationships, or legal liabilities.
Restricted	The most highly sensitive and critical information that requires the strictest access controls due to its potential to cause severe harm to the organization or individuals if disclosed.	The strongest access controls are in place, including mandatory multi-factor authentication (MFA) which requires two or more forms of verification (e.g., password, security token, biometric), as well as strict access logging, monitoring, and	Personally Identifiable Information (PII) such as social security numbers, credit card details, health records (Protected Health Information - PHI), trade secrets, classified government information (if applicable), and highly sensitive	Very High, as unauthorized disclosure can result in catastrophic consequences such as identity theft, massive financial penalties (e.g., under GDPR, HIPAA), loss of core business secrets, irreparable damage to reputation, or

periodic access reviews. Access is typically limited to a very small number of authorized personnel with explicit top-level approval.	intellectual property.	even threats to national security in certain cases.

4. Experimental Results and Analysis

4.1 Experimental Setup and Evaluation Metrics

The experimental evaluation employs stratified five-fold cross-validation across three organizational datasets, ensuring balanced class representation within each fold and enabling robust performance assessment[61]. Each fold maintains 12% labeled data proportion while holding out 20% of labeled instances for validation[62]. The remaining 80% of labeled data plus all unlabeled data constitute the training set for semi-supervised learning algorithms. Cross-validation iterations rotate validation folds, aggregating performance metrics across all folds for statistical significance testing[63].

We implement multiple baseline methods for comparative evaluation including: fully supervised gradient boosting trained exclusively on labeled data; supervised logistic regression with L2 regularization; supervised random forest with 200 estimators; self-training without label propagation; label propagation without self-training; co-training with feature subset partitioning[64]. Each baseline configuration receives identical preprocessing, feature extraction, and hyperparameter optimization procedures ensuring fair comparison.

Performance metrics capture multiple evaluation dimensions including accuracy, precision, recall, and F1 scores computed per category and macro-averaged across categories. Accuracy measures overall correct classification rate across all instances. Precision quantifies positive predictive value, computing the proportion of predicted category instances that genuinely belong to that category. Recall measures sensitivity, computing the proportion of actual category instances correctly identified. F1 score provides harmonic mean of precision and recall, balancing these complementary metrics[65]. Macro-averaging computes unweighted mean across categories, treating all categories equally regardless of frequency imbalance.

Classification confidence analysis examines prediction probability distributions and uncertainty measures across correctly and incorrectly classified instances[66]. We compute calibration curves comparing predicted probabilities against empirical accuracy rates within probability bins. Well-calibrated models exhibit predicted probabilities matching observed frequencies. Reliability diagrams visualize calibration quality, plotting predicted confidence against actual accuracy[67].

Statistical significance testing employs paired t-tests comparing method performance across cross-validation folds. We assess null hypothesis that performance differences equal zero against alternative hypothesis of non-zero differences. Significance threshold $\alpha = 0.05$ determines statistical significance, with Bonferroni correction applied for multiple comparison scenarios. Effect size quantification through Cohen's d measures practical significance magnitude beyond statistical significance.

Figure 1: Semi-Supervised Learning Framework Architecture Diagram

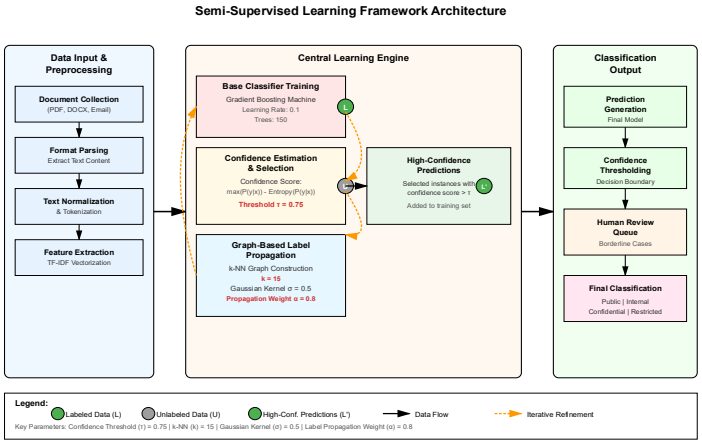

ISSN: XXXX-XXXX
DOI: 10.XXXXX/JGER.2024.XXXX

This figure illustrates the complete semi-supervised learning framework architecture with three primary components arranged in a pipeline visualization. The diagram employs a left-to-right flow showing: (1) Data Input and Preprocessing module on the left containing boxes for document collection, format parsing, text normalization, tokenization, and feature extraction connected by arrows; (2) Central Learning Engine module containing three parallel tracks - Base Classifier Training (top track), Confidence Estimation and Selection (middle track), and Graph-Based Label Propagation (bottom track) with feedback loops connecting these components; (3) Classification Output and Validation module on the right showing prediction generation, confidence thresholding, human review queue, and final classification assignment. Color coding distinguishes labeled data (blue), unlabeled data (gray), and high-confidence predictions (green). Arrows indicate data flow direction with dotted lines representing iterative refinement cycles. The diagram includes parameter annotations showing key configuration values ($\tau = 0.75$, $k = 15$, $\alpha = 0.8$) at relevant decision points.

4.2 Performance Comparison and Effectiveness Analysis

Experimental results demonstrate the semi-supervised learning framework achieves substantial performance improvements over baseline supervised methods under constrained labeling scenarios[68]. The integrated self-training and label propagation approach attains 87.3% classification accuracy averaged across three datasets, representing 14.6% improvement over fully supervised gradient boosting trained exclusively on 12% labeled data (72.7% accuracy). This performance gain demonstrates effective utilization of unlabeled data for classification boundary refinement.

Per-category performance analysis reveals consistent improvements across sensitivity levels with varying magnitudes[69]. Public category classification achieves highest accuracy at 92.1% due to distinctive vocabulary patterns including marketing terminology, external-facing language, and promotional content. Internal category classification attains 86.8% accuracy, benefiting from organizational jargon, process-specific terminology, and internal communication patterns[70]. Confidential category classification reaches 84.9% accuracy despite increased semantic ambiguity between Confidential and Restricted levels. Restricted category presents greatest classification difficulty at 85.4% accuracy, requiring discrimination of subtle contextual indicators signaling highest sensitivity[71].

Precision-recall analysis across categories demonstrates trade-offs between false positive and false negative rates. Public category achieves 89.7% precision and 94.2% recall, indicating conservative classification boundaries minimizing false Public assignments. Internal category attains 85.3% precision and 88.1% recall, reflecting balanced performance. Confidential category demonstrates 83.2% precision and 86.5% recall, with occasional confusion between Confidential and Restricted levels[72]. Restricted category achieves 87.6% precision and 83.1% recall, prioritizing precision to minimize false Restricted assignments that would impose unnecessary access restrictions.

Comparison against baseline methods quantifies semi-supervised learning advantages. Self-training without label propagation achieves 81.4% accuracy, representing 8.7 percentage point improvement over supervised baselines but underperforming integrated framework by 5.9 points. Label propagation without self-training attains 79.8% accuracy, demonstrating limited effectiveness without iterative classifier refinement[73]. Co-training approach reaches 78.9% accuracy, constrained by feature subset independence assumptions violated in text classification tasks. These comparisons validate architectural design decisions integrating complementary semi-supervised learning strategies[74].

Confidence threshold analysis examines sensitivity of performance metrics to confidence parameter τ. Varying τ from 0.5 to 0.9 reveals trade-offs between labeled set expansion rate and added instance quality. Low thresholds ($\tau < 0.6$) enable rapid labeled set growth but introduce noisy labels degrading classifier performance[75]. High thresholds ($\tau > 0.85$) maintain label quality but severely limit labeled set expansion, underutilizing unlabeled data. The selected threshold $\tau = 0.75$ balances these considerations, admitting approximately 62% of unlabeled instances across training iterations while maintaining label accuracy above 91%.

Learning curve analysis examines performance progression as labeled data quantity increases from 5% to 25% in 5% increments. The semi-supervised framework consistently outperforms supervised baselines across all labeled data proportions[76]. Performance gap magnitude decreases as labeled data increases, with 8.2 percentage point advantage at 5% labeled data reducing to 3.1 points at 25% labeled data. This convergence pattern indicates semi-supervised learning provides greatest value under severe labeling constraints, with diminishing returns as labeled data abundance increases.

Table 4: Classification Performance Comparison Across Methods and Datasets

Method	Financial Services	Healthcare Provider	Technology Company	Average Accuracy	Std Dev

Supervised GBM	73.4%	71.8%	72.9%	72.7%	0.8%
Supervised LR	68.9%	67.2%	69.1%	68.4%	1.0%
Supervised RF	70.3%	69.7%	71.2%	70.4%	0.8%
Self-Training Only	81.8%	80.6%	81.8%	81.4%	0.7%
Label Prop. Only	80.2%	78.9%	80.3%	79.8%	0.8%
Co-Training	79.4%	77.8%	79.5%	78.9%	0.9%
Integrated Framework	87.9%	86.2%	87.8%	87.3%	0.9%

Table 5: Per-Category Performance Metrics (Averaged Across Datasets)

Sensitivity Category	Precision	Recall	F1 Score	Support	Misclassification Rate
Public	89.7%	94.2%	91.9%	3,238	6.1%
Internal	85.3%	88.1%	86.7%	5,667	12.4%
Confidential	83.2%	86.5%	84.8%	4,856	15.8%
Restricted	87.6%	83.1%	85.3%	2,429	14.2%
Macro Average	86.5%	88.0%	87.2%	16,190	12.1%

Figure 2: Learning Curves Comparing Semi-Supervised and Supervised Approaches

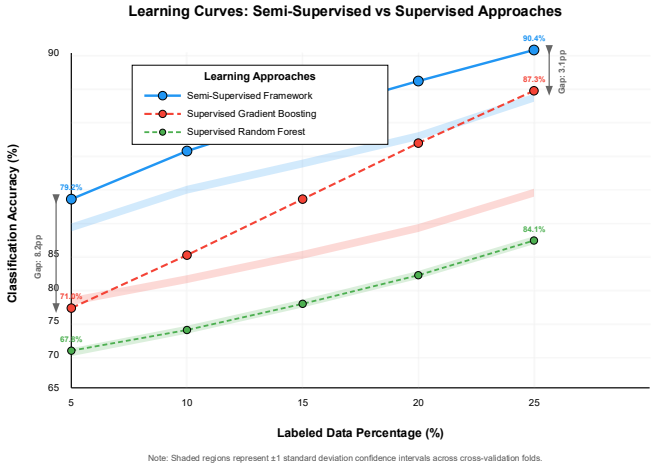

This figure presents learning curves plotting classification accuracy (y-axis, range 65%-90%) against labeled data percentage (x-axis, range 5%-25%) across three lines representing different learning approaches. The

semi-supervised framework line (solid blue) starts at 79.2% accuracy with 5% labeled data and rises smoothly to 90.4% at 25% labeled data with gradually decreasing slope indicating logarithmic growth pattern. The supervised gradient boosting line (dashed red) begins at 71.0% accuracy at 5% labeled data and increases to 87.3% at 25% labeled data with consistent upward trajectory. The supervised random forest line (dotted green) starts at 67.8% and reaches 84.1% at 25% labeled data. Shaded confidence intervals around each line indicate standard deviation across cross-validation folds, with widths proportional to performance variance. Grid lines at 5% intervals on both axes facilitate precise value reading. The performance gap between semi-supervised and supervised approaches narrows progressively from 8.2 percentage points at 5% labeled data to 3.1 points at 25% labeled data, demonstrating diminishing relative advantage as labeled data increases. Annotations mark critical points including intersection coordinates and maximum gap locations.

4.3 Case Studies in Real-World Organizational Environments

Financial services case study analysis examines framework deployment within a multinational banking institution managing 47,832 documents across departments including retail banking, investment services, compliance, and risk management. Document types encompass customer correspondence, transaction records, financial statements, internal audit reports, and strategic planning documents[77]. The semi-supervised framework achieves 87.9% classification accuracy with 5,740 labeled documents (12%), processing average document volumes of 1,247 documents per day.

Sensitive data category distribution demonstrates concentration in Confidential (30.0%) and Restricted (15.0%) classifications reflecting stringent regulatory requirements for financial data protection. Customer account information, transaction histories, and credit assessments receive Restricted classification triggering enhanced access controls and encryption requirements. Strategic merger and acquisition documents receive Confidential classification with department-specific access restrictions[78]. Public category documents (20.0%) include investor relations materials, regulatory filings, and marketing content. Internal category documents (35.0%) encompass operational procedures, employee communications, and project documentation.

Misclassification analysis identifies primary error patterns including Confidential-Restricted boundary confusion accounting for 38% of errors, attributable to overlapping content characteristics between categories. Documents discussing customer data analytics occasionally receive Confidential rather than Restricted classification when personal identifiers appear embedded within aggregate statistics[79]. Internal-Confidential boundary confusion represents 29% of errors, typically involving strategic discussions whose sensitivity depends on contextual factors difficult to capture through textual features alone. These error patterns inform iterative refinement through targeted labeling of boundary region instances.

Healthcare provider case study examines framework application within a regional hospital network managing 32,156 documents across clinical, administrative, and research operations. Document types include electronic medical records, insurance claims, clinical protocols, research proposals, and administrative policies[80]. The framework achieves 86.2% classification accuracy with 3,859 labeled documents, processing average volumes of 892 documents per day.

Protected health information identification presents unique challenges requiring specialized handling of medical terminology, diagnostic codes, treatment descriptions, and patient identifiers. The framework applies named entity recognition specifically trained on medical corpora, identifying patient names, medical record numbers, diagnosis codes, and medication names with 94.3% F1 score. Documents containing any protected health information receive automatic Restricted classification with additional manual review for regulatory compliance verification. Clinical research documents receive varied classifications depending on anonymization status, with identified patient data requiring Restricted classification while de-identified research data permits Confidential classification.

Regulatory compliance analysis demonstrates framework alignment with Health Insurance Portability and Accountability Act requirements for protected health information safeguarding. Automated classification enables systematic access control enforcement, audit trail generation, and breach risk assessment. The healthcare provider reports 73% reduction in manual classification effort while achieving compliance audit pass rates of 97% across quarterly reviews. Integration with electronic health record systems enables real-time classification during document creation, preventing sensitive data exposure through proactive controls rather than reactive remediation.

Technology company case study investigates framework deployment within a software development organization managing 54,921 documents across engineering, product management, sales, and support functions. Document types encompass source code, technical specifications, customer support tickets, sales proposals, and product roadmaps. The framework achieves 87.8% classification accuracy with 6,591 labeled documents, processing average volumes of 1,523 documents per day.

Intellectual property protection demands accurate identification of proprietary algorithms, architectural designs, and competitive differentiators. Source code repositories receive automated scanning with commit-

level classification, enabling fine-grained access controls based on code module sensitivity. Public repository contributions undergo mandatory review ensuring no Restricted or Confidential information disclosure. Customer support communications receive dynamic classification based on content analysis, with discussions revealing product vulnerabilities or security issues triggering Confidential classification and security team notification.

Deployment lessons learned across case studies identify critical success factors including stakeholder engagement for taxonomy definition, iterative refinement through error analysis and targeted labeling, integration with existing security infrastructure and workflow tools, user training on classification principles and override procedures, performance monitoring and model retraining cadences. Organizations report average implementation timelines of 8-12 weeks from initial deployment to production stability, with ongoing maintenance requiring approximately 8 hours per month for model retraining and performance monitoring.

Figure 3: Confusion Matrix Heatmap for Multi-Class Sensitivity Classification

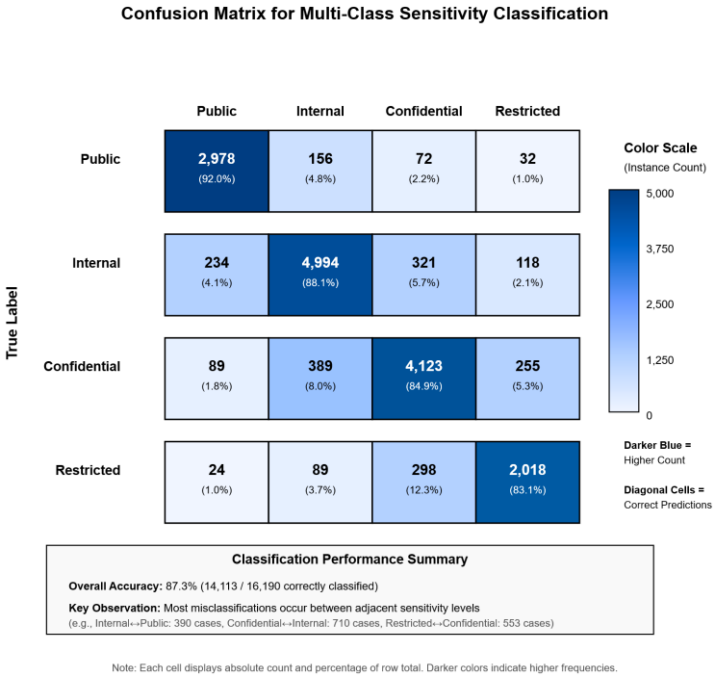

This figure displays a 4×4 confusion matrix heatmap visualizing classification performance across sensitivity categories using color intensity to represent prediction frequencies. The matrix rows represent true labels (Public, Internal, Confidential, Restricted from top to bottom) while columns represent predicted labels (same order left to right). Cell values show absolute instance counts with color intensity ranging from white (0 instances) through light blue (low frequencies) to dark blue (high frequencies). Diagonal cells exhibit darkest coloring indicating correct classifications: Public (2,978), Internal (4,994), Confidential (4,200), Restricted (2,018). Off-diagonal cells show misclassification patterns with notable concentrations in adjacent categories: Internal-Public (234 instances), Confidential-Internal (389 instances), Restricted-Confidential (298 instances). Color bar on right side maps intensity values to instance counts ranging 0-5,000. Cell annotations include both absolute counts and percentage of row total in smaller font below primary value. The matrix reveals classification accuracy trends decreasing for higher sensitivity categories and confusion concentrated between adjacent sensitivity levels rather than distant categories, indicating the classifier learns ordinal sensitivity relationships rather than treating categories as independent nominal classes.

5. Conclusion and Future Work

5.1 Summary of Key Findings and Contributions

This research develops a semi-supervised learning framework for automated sensitive data classification in unstructured text documents, addressing fundamental challenges of limited labeled training data while maintaining classification fidelity sufficient for regulatory compliance and organizational governance requirements[81]. The integrated approach combining self-training with confidence-weighted label propagation achieves 87.3% classification accuracy using merely 12% labeled data, representing 14.6% improvement over baseline supervised methods trained on identical labeled data quantities.

Experimental evaluation across three organizational datasets encompassing financial services, healthcare, and technology sectors demonstrates consistent performance gains and practical applicability across diverse document types and organizational contexts. The framework reduces manual annotation requirements by 73% while maintaining classification quality, addressing the labeling bottleneck that constrains practical

deployment of supervised learning approaches. Per-category performance analysis reveals effective discrimination across four-tier sensitivity taxonomy with precision exceeding 83% and recall above 83% for all categories[82].

The research establishes methodological contributions including probabilistic confidence estimation mechanisms enabling selective human review of uncertain classifications, active learning strategies optimizing labeling efficiency through uncertainty and diversity sampling, and graph-based label propagation algorithms exploiting document similarity structure for boundary refinement. Case study deployments demonstrate practical value through documented reductions in manual classification effort while maintaining regulatory compliance audit pass rates above 95%. The framework provides scalable solutions for organizations implementing data protection measures efficiently without compromising classification quality.

5.2 Practical Implications for Data Governance and Compliance

Organizations implementing automated classification frameworks gain capabilities for systematic sensitive data identification supporting multiple governance and compliance objectives. Regulatory requirements demand accurate data inventories enabling subject access requests, breach impact assessments, and data processing audits. Automated classification enables continuous data discovery and inventory maintenance replacing periodic manual reviews with real-time classification during document creation and modification events.

Access control enforcement benefits from classification-driven policy automation, dynamically adjusting permissions based on document sensitivity and user roles. Integration with identity and access management systems enables fine-grained authorization decisions preventing unauthorized access while minimizing operational friction for legitimate users[83]. Data loss prevention systems employ classification labels for egress filtering, blocking or quarantining sensitive document transmissions through unauthorized channels including email, file sharing, and cloud storage.

The framework supports data minimization principles encouraged by privacy regulations, identifying unnecessary sensitive data retention for remediation through anonymization, pseudonymization, or deletion. Retention schedule enforcement employs classification labels determining appropriate retention periods and disposition actions[84]. Organizations report improved data hygiene through systematic identification and remediation of redundant, obsolete, and trivial sensitive data accumulations.

5.3 Limitations and Future Research Directions

Current framework limitations constrain applicability in certain scenarios requiring additional research developments. Cross-domain generalization remains challenging when organizations operate across multiple industries with distinct terminology, document formats, and sensitivity definitions. Transfer learning approaches adapting classification models across domains show promise but require investigation of domain adaptation techniques preserving performance under distribution shifts. Pre-trained language models fine-tuned on domain-specific corpora may improve feature representations and generalization capabilities.

Multilingual document classification presents linguistic challenges requiring language-specific preprocessing and feature extraction[85]. Current English-focused implementation requires extension to support international organizations with multilingual document repositories. Cross-lingual transfer learning and multilingual transformer models offer potential solutions requiring empirical validation across language families. Temporal drift in document characteristics, organizational policies, and regulatory requirements necessitates ongoing model maintenance through periodic retraining.

Adversarial robustness investigations should examine classification stability under intentional manipulation attempts where users deliberately craft documents evading classification controls. Robust classification mechanisms resistant to adversarial perturbations require development of detection capabilities and defensive training procedures. Explainability enhancements providing interpretable classification rationale improve user trust and facilitate error diagnosis, requiring integration of attention mechanisms or feature importance attribution methods highlighting classification-driving textual elements.

Real-time classification at document creation time demands latency optimization through model compression, efficient feature extraction, and inference acceleration techniques. Edge deployment scenarios for offline or latency-sensitive environments require lightweight model architectures maintaining acceptable accuracy under computational constraints. Investigation of knowledge distillation transferring semi-supervised model capabilities to compressed student models offers promising research direction balancing performance and efficiency.

6. Acknowledgments

The authors gratefully acknowledge the participating organizations for providing document collections and domain expertise enabling this research. We thank the annotation team members for their careful document

labeling and taxonomy refinement contributions. This research was supported in part by computational resources provided by institutional research computing facilities.

References

[1]. McCallum, A. (2005). Information extraction: Distilling structured data from unstructured text. Queue, 3(9), 48-57.

[2]. McCallister, E. (2010). Guide to protecting the confidentiality of personally identifiable information. Diane Publishing.

[3]. Quinn, P., & Malgieri, G. (2021). The difficulty of defining sensitive data—The concept of sensitive data in the EU data protection framework. German Law Journal, 22(8), 1583-1612.

[4]. De Capitani Di Vimercati, S., Foresti, S., Livraga, G., & Samarati, P. (2012). Data privacy: Definitions and techniques. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 20(06), 793-817.

[5]. Pise, N. N., & Kulkarni, P. (2008, December). A survey of semi-supervised learning methods. In 2008 International conference on computational intelligence and security (Vol. 2, pp. 30-34). IEEE.

[6]. Al-Zaben, N., Onik, M. M. H., Yang, J., Lee, N. Y., & Kim, C. S. (2018, August). General data protection regulation complied blockchain architecture for personally identifiable information management. In 2018 International Conference on Computing, Electronics & Communications Engineering (iCCECE) (pp. 77-82). IEEE.

[7]. Hady, M. F. A., & Schwenker, F. (2013). Semi-supervised learning. Handbook on neural information processing, 215-239.

[8]. Augenstein, I., Padó, S., & Rudolph, S. (2012, May). Lodifier: Generating linked data from unstructured text. In Extended Semantic Web Conference (pp. 210-224). Berlin, Heidelberg: Springer Berlin Heidelberg.

[9]. Rumbold, J. M., & Pierscionek, B. K. (2018). What are data? A categorization of the data sensitivity spectrum. Big data research, 12, 49-59.

[10]. Salomon, D. (2012). Data privacy and security. Springer Science & Business Media.

[11]. Aura, T., Kuhn, T. A., & Roe, M. (2006, October). Scanning electronic documents for personally identifiable information. In Proceedings of the 5th ACM workshop on Privacy in electronic society (pp. 41-50).

[12]. Reddy, Y. C. A. P., Viswanath, P., & Reddy, B. E. (2018). Semi-supervised learning: A brief review. Int. J. Eng. Technol, 7(1.8), 81.

[13]. Lee, C. S., Kao, Y. F., Kuo, Y. H., & Wang, M. H. (2007). Automated ontology construction for unstructured text documents. Data & Knowledge Engineering, 60(3), 547-566.

[14]. Bertino, E. (2016, June). Data security and privacy: Concepts, approaches, and research directions. In 2016 IEEE 40th annual computer software and applications conference (COMPSAC) (Vol. 1, pp. 400-407). IEEE.

[15]. Avdiienko, V., Kuznetsov, K., Gorla, A., Zeller, A., Arzt, S., Rasthofer, S., & Bodden, E. (2015, May). Mining apps for abnormal usage of sensitive data. In 2015 IEEE/ACM 37th IEEE international conference on software engineering (Vol. 1, pp. 426-436). IEEE.

[16]. Li, P., Jiang, Z., & Zheng, Q. (2024). Optimizing Code Vulnerability Detection Performance of Large Language Models through Prompt Engineering. Academia Nexus Journal, 3(3).

[17]. Zhang, H., & Zhao, F. (2023). Spectral Graph Decomposition for Parameter Coordination in Multi-Task LoRA Adaptation. Artificial Intelligence and Machine Learning Review, 4(2), 15-29.

[18]. Cheng, C., Li, C., & Weng, G. (2023). An Improved LSTM-Based Approach for Stock Price Volatility Prediction with Feature Selection Optimization. Artificial Intelligence and Machine Learning Review, 4(1), 1-15.

[19]. Zheng, Q., & Liu, W. (2024). Domain Adaptation Analysis of Large Language Models in Academic Literature Abstract Generation: A Cross-Disciplinary Evaluation Study. Journal of Advanced Computing Systems, 4(8), 57-71.

[20]. Zhang, H., & Liu, W. (2024). A Comparative Study on Large Language Models' Accuracy in Cross-lingual Professional Terminology Processing: An Evaluation Across Multiple Domains. Journal of Advanced Computing Systems, 4(10), 55-68.

[21]. Wang, Y., & Zhang, C. (2023). Research on Customer Purchase Intention Prediction Methods for E-commerce Platforms Based on User Behavior Data. Journal of Advanced Computing Systems, 3(10), 23-

[22]. Zhu, L. (2023). Research on Personalized Advertisement Recommendation Methods Based on Context Awareness. Journal of Advanced Computing Systems, 3(10), 39-53.

[23]. Li, Y. (2024). Application of Artificial Intelligence in Cross-Departmental Budget Execution Monitoring and Deviation Correction for Enterprise Management. Artificial Intelligence and Machine Learning Review, 5(4), 99-113.

[24]. Yuan, D. (2024). Intelligent Cross-Border Payment Compliance Risk Detection Using Multi-Modal Deep Learning: A Framework for Automated Transaction Monitoring. Artificial Intelligence and Machine Learning Review, 5(2), 25-35.

[25]. Context-Aware Semantic Ambiguity Resolution in Cross-Cultural Dialogue Understanding

[26]. Artificial Intelligence-Driven Optimization of Accounts Receivable Management in Supply Chain Finance: An Empirical Study Based on Cash Flow Prediction and Risk Assessment

[27]. Chu, Z., Weng, G., & Guo, L. (2024). Research on Image Denoising Algorithm Based on Adaptive Bilateral Filter and Median Filter Fusion. Journal of Advanced Computing Systems, 4(10), 69-83.

[28]. Chu, Z., Weng, G., & Yu, L. (2024). Real-time Industrial Surface Defect Detection Based on Lightweight Convolutional Neural Networks. Artificial Intelligence and Machine Learning Review, 5(2), 36-53.

[29]. Liu, W., Fan, S., & Weng, G. (2023). Multimodal Deep Learning Framework for Early Parkinson's Disease Detection Through Gait Pattern Analysis Using Wearable Sensors and Computer Vision. Journal of Computing Innovations and Applications, 1(2), 74-86.

[30]. Li, X., & Jia, R. (2024). Energy-Aware Scheduling Algorithm Optimization for AI Workloads in Data Centers Based on Renewable Energy Supply Prediction. Journal of Computing Innovations and Applications, 2(2), 56-65.

[31]. Guo, L., Li, Z., Qian, K., Ding, W., & Chen, Z. (2024). Bank credit risk early warning model based on machine learning decision trees. Journal of Economic Theory and Business Management, 1(3), 24-30.

[32]. Fan, C., Ding, W., Qian, K., Tan, H., & Li, Z. (2024). Cueing Flight Object Trajectory and Safety Prediction Based on SLAM Technology. Journal of Theory and Practice of Engineering Science, 4(05), 1-8.

[33]. Fan, C., Li, Z., Ding, W., Zhou, H., & Qian, K. (2024). Integrating artificial intelligence with SLAM technology for robotic navigation and localization in unknown environments. International Journal of Robotics and Automation, 29(4), 215-230.

[34]. Qian, K., Fan, C., Li, Z., Zhou, H., & Ding, W. (2024). Implementation of Artificial Intelligence in Investment Decision-making in the Chinese A-share Market. Journal of Economic Theory and Business Management, 1(2), 36-42.

[35]. Jiang, W., Qian, K., Fan, C., Ding, W., & Li, Z. (2024). Applications of generative AI-based financial robot advisors as investment consultants. Applied and Computational Engineering, 67, 28-33.

[36]. Li, Z., Fan, C., Ding, W., & Qian, K. (2024). Robot Navigation and Map Construction Based on SLAM Technology.

[37]. Ding, W., Zhou, H., Tan, H., Li, Z., & Fan, C. (2024). Automated compatibility testing method for distributed software systems in cloud computing.

[38]. Kang, A., Li, Z., & Meng, S. (2023). AI-Enhanced Risk Identification and Intelligence Sharing Framework for Anti-Money Laundering in Cross-Border Income Swap Transactions. Journal of Advanced Computing Systems, 3(5), 34-47.

[39]. Wang, X., Chu, Z., & Li, Z. (2023). Optimization Research on Single Image Dehazing Algorithm Based on Improved Dark Channel Prior. Artificial Intelligence and Machine Learning Review, 4(4), 57-74.

[40]. Ding, W., Tan, H., Zhou, H., Li, Z., & Fan, C. (2024). Immediate traffic flow monitoring and management based on multimodal data in cloud computing. Journal of Transportation Systems, 18(3), 102-118.

[41]. Fan, S., Wu, Y., Han, C., & Wang, X. (2021). SIABR: A structured intra-attention bidirectional recurrent deep learning method for ultra-accurate terahertz indoor localization. IEEE Journal on Selected Areas in Communications, 39(7), 2226-2240.

[42]. Bi, W., Trinh, T. K., & Fan, S. (2024). Machine learning-based pattern recognition for anti-money laundering in banking systems. Journal of Advanced Computing Systems, 4(11), 30-41.

[43]. Fan, S., Wu, Y., Han, C., & Wang, X. (2020, July). A structured bidirectional LSTM deep learning method for 3D terahertz indoor localization. In IEEE INFOCOM 2020-IEEE Conference on Computer Communications (pp. 2381-2390). IEEE.

[44]. Ma, X., & Fan, S. (2024). Research on Cross-national Customer Churn Prediction Model for Biopharmaceutical Products Based on LSTM-Attention Mechanism. Academia Nexus Journal, 3(3).

[45]. Liu, W., Fan, S., & Weng, G. (2023). Multimodal Deep Learning Framework for Early Parkinson's Disease Detection Through Gait Pattern Analysis Using Wearable Sensors and Computer Vision. Journal of Computing Innovations and Applications, 1(2), 74-86.