GEORGE BROWN PRESS
JOURNALS, BOOKS & PROCEEDINGS

Journal of Computer, Signal, and System Research

Article

Privacy-Preserving AI in SMB Customer Service: Balancing Data Isolation, Compliance, and Automation Efficacy

Zhenyuan He [1,*]

[1] Walmart Global Tech, Sunnyvale, CA, USA
* Correspondence: Zhenyuan He, Walmart Global Tech, Sunnyvale, CA, USA

Abstract: This research investigates the application of Privacy-Preserving AI (PPAI) techniques in Small and Medium-sized Business (SMB) customer service environments, recognizing that SMBs constitute the backbone of the U.S. economy. It focuses on balancing the seemingly conflicting demands of data isolation, regulatory compliance (e.g., GDPR, CCPA), and the efficacy of AI-driven automation. The study explores various PPAI methodologies, including Federated Learning, Differential Privacy, Homomorphic Encryption, and Secure Multi-Party Computation, assessing their suitability for different SMB customer service scenarios. We analyze the trade-offs between privacy guarantees, model accuracy, computational overhead, and implementation complexity. Real-world case studies and simulations are used to evaluate the performance of selected PPAI techniques across key customer service metrics such as response time, customer satisfaction, and issue resolution rate. Furthermore, the research addresses the practical challenges of deploying PPAI in resource-constrained SMBs, considering factors like data heterogeneity, limited technical expertise, and cost considerations. This study proposes a framework for SMBs to effectively adopt PPAI in their customer service operations, ensuring robust data protection, adherence to compliance regulations, and enhanced automation capabilities. By lowering the technical barrier for SMBs to adopt AI while strictly adhering to privacy laws, this work aims to support widespread economic modernization alongside robust consumer privacy protection. The findings provide valuable insights for SMBs, AI developers, and policymakers aiming to promote the responsible and ethical use of AI in customer service.

Keywords: privacy-preserving AI; federated learning; differential privacy; homomorphic encryption; customer service automation; SMB; GDPR; CCPA

Received: 25 December 2025
Revised: 05 February 2026
Accepted: 16 February 2026
Published: 21 February 2026

1. Introduction

1.1. Background and Motivation

The integration of Artificial Intelligence (AI) into customer service operations is rapidly transforming how businesses interact with their clientele. AI-powered chatbots, personalized recommendations, and automated support systems are becoming increasingly prevalent, promising enhanced efficiency and improved customer satisfaction. However, this technological advancement coincides with growing concerns surrounding data privacy and stringent regulatory frameworks like GDPR and CCPA [1]. Small and medium-sized businesses (SMBs), which form the backbone of the U.S. economy, face unique challenges in navigating this complex landscape. They often lack the resources and expertise to implement robust privacy-preserving mechanisms while simultaneously leveraging the full potential of AI for customer service automation. The core issue lies in balancing the need for data utility, represented by the variable U, with the imperative of data protection, quantified as P, within the resource constraints R typical of SMBs. Enabling SMBs to adopt AI securely and compliantly is thus not only a technical

challenge but also a matter of supporting broader economic competitiveness and modernization [2].

1.2. Problem Statement and Research Objectives

The increasing reliance on AI in Small and Medium-sized Business (SMB) customer service presents a critical challenge: how to leverage AI's automation capabilities while adhering to stringent data privacy regulations like GDPR and maintaining data isolation. Many SMBs lack the resources to implement sophisticated privacy-preserving techniques, hindering their ability to fully utilize AI. This research investigates methods for achieving effective AI-driven customer service in SMBs without compromising customer data privacy [3].

The primary objective is to identify and evaluate practical privacy-preserving AI techniques suitable for SMB customer service contexts. Specifically, this study aims to: 1) assess the feasibility of federated learning (FL) for training AI models on decentralized customer data; 2) explore the effectiveness of differential privacy (DP) in safeguarding sensitive information during model deployment; and 3) develop a framework that balances privacy guarantees, model accuracy, and computational costs for SMBs [4]. Key research questions include: How can FL be adapted to address the unique data heterogeneity challenges in SMB customer service? What are the optimal DP parameters to minimize privacy loss while preserving model utility? And, what is the total cost of ownership for implementing privacy-preserving AI solutions in SMBs?

2. Literature Review

2.1. Privacy-Preserving AI Techniques

Privacy-Preserving AI (PPAI) techniques are crucial for enabling AI applications while safeguarding sensitive data, particularly relevant in Small and Medium-sized Businesses (SMBs) dealing with customer information. Federated Learning (FL) allows model training across decentralized devices or servers holding local data samples, without directly exchanging the data itself. Instead, only model updates are shared, as explored by researchers in various domains. Differential Privacy (DP) adds carefully calibrated noise to the data or model outputs, ensuring that the presence or absence of any single data point has a limited impact on the result. The parameter ϵ controls the privacy level; a smaller ϵ provides stronger privacy but potentially lower utility [5]. Homomorphic Encryption (HE) enables computations on encrypted data, meaning that AI models can be trained and used without ever decrypting the underlying information. Different HE schemes offer varying levels of security and computational efficiency. Secure Multi-Party Computation (SMPC) allows multiple parties to jointly compute a function over their private inputs, without revealing those inputs to each other. This is particularly useful when data is distributed across multiple SMBs. Other techniques, such as k-anonymity and data masking, also contribute to privacy preservation, although they may offer weaker privacy guarantees compared to FL, DP, HE, and SMPC. The choice of PPAI technique depends on the specific application, data characteristics, and desired trade-off between privacy and utility, as shown in Figure 1.

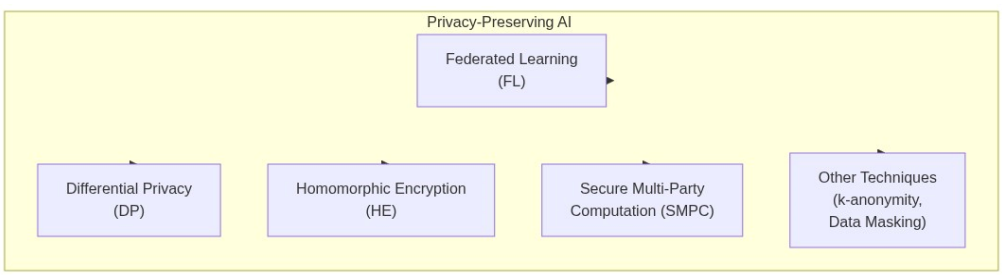

Figure 1. Taxonomy of Privacy-Preserving AI Techniques.

2.2. AI in Customer Service Applications

AI is increasingly prevalent in customer service, automating tasks and enhancing customer experiences. Chatbots, powered by natural language processing (NLP), handle routine inquiries, freeing human agents for complex issues. Sentiment analysis, another key application, uses machine learning to gauge customer emotions from text and speech, enabling proactive intervention and personalized responses [6]. Furthermore, AI algorithms drive personalized recommendations, suggesting products or services tailored to individual customer preferences, thereby increasing sales and satisfaction. These applications leverage techniques like machine learning, deep learning, and NLP to improve efficiency and effectiveness in customer service operations [7]. The variable x can represent customer satisfaction, and y can represent sales increase.

2.3. Challenges and Opportunities for SMBs

SMBs face distinct hurdles in adopting Privacy-Preserving AI (PPAI) for customer service. Limited budgets restrict investment in sophisticated PPAI technologies and specialized personnel. Data scarcity, with smaller customer datasets compared to larger enterprises, can hinder the training and effectiveness of PPAI models. However, opportunities exist. SMBs' agility allows for quicker experimentation and implementation of novel PPAI solutions [8]. A focus on specific customer segments enables the development of tailored PPAI models, potentially achieving higher accuracy with less data. Furthermore, PPAI can be a competitive differentiator, attracting privacy-conscious customers and building trust, thereby increasing customer lifetime value (LTV). As summarized in Table 1, different AI solutions exhibit distinct trade-offs in cost, privacy protection, and performance for SMB customer service.

Table 1. Comparative Analysis of AI Solutions for SMB Customer Service.

Feature	Traditional AI	Privacy-Preserving AI (PPAI)
Cost	Generally lower initial investment	Higher initial investment due to specialized technologies and expertise
Data Requirements	Requires substantial amounts of data for effective training	Can be effective with smaller, more targeted datasets, especially for tailored models designed for specific customer segments
Data Privacy	Potentially exposes sensitive customer data during training and deployment	Protects customer data privacy through techniques like federated learning, differential privacy, and homomorphic encryption
Training Data Access	Requires direct access to raw customer data	Can train models on decentralized or anonymized data, reducing the need for direct access to raw customer data
Model Security	Vulnerable to data breaches and privacy violations	Enhanced security features to prevent data leakage and protect sensitive information
Compliance	Requires rigorous data governance and compliance measures to meet privacy regulations (e.g., GDPR, CCPA)	Facilitates compliance with privacy regulations by design, reducing the risk of penalties
Customer Trust	Potential erosion of trust due to privacy concerns	Builds trust with privacy-conscious customers, leading to increased customer loyalty and LTV

Implementation Speed	Potentially faster initial implementation due to readily available solutions	Can take longer initially due to the complexity of PPAI techniques and the need for specialized expertise; SMB agility allows for quicker experimentation, however
Competitive Advantage	Limited differentiation based on AI technology alone	Offers a strong competitive advantage by appealing to privacy concerns and building a reputation for responsible data handling

3. Materials and Methods

3.1. Data Collection and Preprocessing

To evaluate the proposed privacy-preserving AI framework, we utilized a combination of real-world data characteristics and synthetic data generation techniques. Due to the sensitive nature of customer service interactions and the difficulty in obtaining sufficiently large, anonymized datasets from small and medium-sized businesses (SMBs), we opted for a predominantly synthetic approach [9]. This allowed us to control data characteristics, simulate various customer interaction scenarios, and rigorously test the framework's performance under different conditions while adhering to strict privacy constraints.

The synthetic datasets were generated to mimic typical customer service interactions across several SMB sectors, including retail, hospitality, and basic IT support. We modeled customer inquiries, agent responses, and associated metadata such as timestamps, interaction channels (e.g., chat, email), and customer sentiment scores. The generation process involved defining probability distributions for key variables [10]. For instance, the length of a customer inquiry, denoted as L_i, was modeled using a log-normal distribution, $L_i \sim LogNormal(\mu, \sigma)$, where μ and σ were adjusted based on the simulated sector. Similarly, customer sentiment, represented by S_c, was generated using a beta distribution, $S_c \sim Beta(\alpha, \beta)$, allowing us to simulate a range of customer satisfaction levels.

Preprocessing steps included tokenization, stemming, and the removal of stop words. Textual data was transformed into numerical representations using TF-IDF (Term Frequency-Inverse Document Frequency) weighting. Categorical features, such as interaction channel, were one-hot encoded. To simulate data breaches or privacy violations, we introduced controlled levels of noise and adversarial examples into the datasets. This allowed us to assess the robustness of our privacy-preserving mechanisms against potential attacks. Finally, the datasets were partitioned into training, validation, and testing sets with a ratio of 70:15:15, respectively, ensuring a fair evaluation of the model's generalization capabilities. A summary of the synthetic datasets and their key characteristics is provided in Table 2.

Table 2. Dataset Summary.

Feature	Description	Distribution/Encoding	Purpose
Customer Inquiry Length (L_i)	Length of customer's question/request	$L_i \sim$ LogNormal(μ, σ), where μ and σ are sector-dependent	Simulate inquiry complexity
Customer Sentiment (S_c)	Customer's satisfaction level	$S_c \sim$ Beta(α, β)	Simulate customer satisfaction
Agent Response	Textual response from customer service agent	Tokenization, stemming, stop word removal, TF-IDF	Model agent behavior

Timestamp	Time of interaction	Not explicitly stated, but likely a numerical representation	Track interaction timeline
Interaction Channel	Method of communication (e.g., chat, email)	One-hot encoded	Categorize interaction type
Metadata	Additional data related to the interaction	Varies depending on the specific metadata.	Provide context to the interaction.
Noise/Adversarial Examples	Introduced to simulate data breaches/attacks	Controlled levels of noise and adversarial examples	Evaluate privacy mechanism robustness

3.2. PPAI Model Implementation and Configuration

The implementation of our privacy-preserving AI (PPAI) model for SMB customer service involved a combination of federated learning (FL) and differential privacy (DP) techniques. The core model architecture is based on a transformer network, specifically a pre-trained BERT model fine-tuned for sentiment analysis and intent recognition on customer service interactions. This choice was motivated by BERT's strong performance in natural language understanding tasks and its adaptability to various downstream applications.

For federated learning, we adopted a client-server architecture, as shown in the system architecture in Figure 2. Each SMB acts as a client, possessing its own local dataset of customer interactions. The central server orchestrates the training process without directly accessing the clients' data. In each training round, the server selects a subset of clients to participate. These clients download the current global model from the server, fine-tune it on their local data, and then upload the updated model weights back to the server. The server aggregates these updates using a weighted averaging approach, where the weights are proportional to the size of each client's dataset [11]. This aggregated model then becomes the new global model for the next training round. The number of clients participating in each round, denoted as C, was a key hyperparameter.

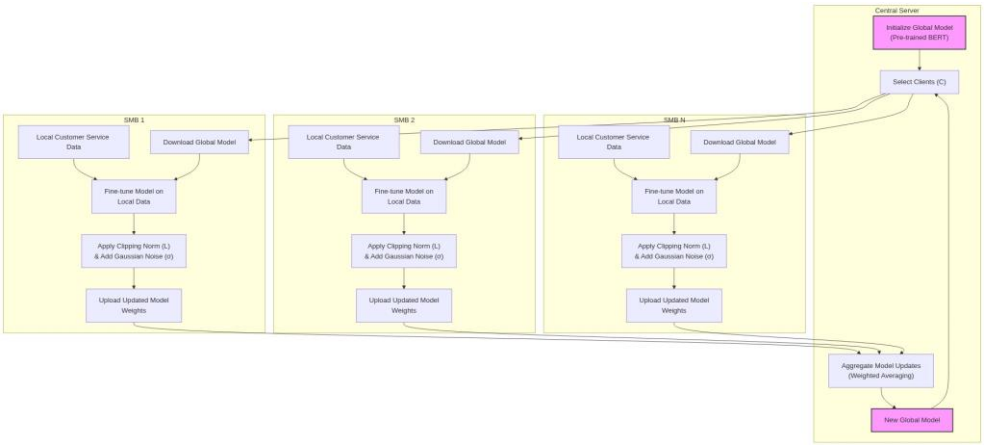

Figure 2. Federated Learning System Architecture for Customer Service Data.

To ensure differential privacy, we implemented Gaussian noise addition to the model updates before they are sent to the server. Specifically, we applied a clipping norm L to the gradients of each client's update to limit the sensitivity of the model to individual data points. Then, we added Gaussian noise with a standard deviation proportional to L and inversely proportional to the privacy parameter ε. The noise scale was determined by

$\sigma = L/\varepsilon$. A smaller ε provides stronger privacy guarantees but can potentially reduce model accuracy. The clipping norm L and the privacy parameter ε were critical hyperparameters that required careful tuning.

The training procedure involved several steps. First, the pre-trained BERT model was initialized with weights from the Hugging Face Transformers library. Second, the model was fine-tuned using the federated learning framework with differential privacy. The training data was preprocessed to remove personally identifiable information (PII) and tokenized using the BERT tokenizer. We used a cross-entropy loss function for both sentiment analysis and intent recognition tasks. The AdamW optimizer was used with a learning rate of $1e-5$.

Hyperparameter tuning was performed using a combination of grid search and random search. We focused on optimizing the number of training rounds, the number of clients participating in each round C, the clipping norm L, the privacy parameter ε, and the learning rate. We evaluated the performance of different hyperparameter configurations based on the accuracy and F1-score on a held-out validation dataset [12]. The validation dataset was constructed to be representative of the overall customer service interaction data across all SMBs, while ensuring no data overlap with the training sets of individual clients. We also monitored the privacy loss using the moments accountant method to ensure that the privacy budget was not exceeded. The optimal hyperparameter settings were selected based on the trade-off between model accuracy and privacy guarantees.

3.3. Evaluation Metrics and Experimental Setup

To comprehensively evaluate the performance of our Privacy-Preserving AI (PPAI) models for SMB customer service, we employed a multifaceted approach, considering accuracy, privacy loss, and computational overhead. Accuracy was measured using standard metrics relevant to the specific task. For classification tasks, such as intent recognition, we used precision, recall, F1-score, and overall accuracy. For natural language generation tasks, like chatbot responses, we employed BLEU score and ROUGE scores to assess the fluency and relevance of the generated text compared to human-generated reference responses.

Privacy loss was quantified using differential privacy metrics. Specifically, we calculated the ϵ and δ values achieved by our differentially private training algorithms. A lower ϵ value indicates a stronger privacy guarantee. We also measured the membership inference attack (MIA) success rate. A lower MIA success rate signifies better protection against adversaries attempting to determine if a specific data point was used in training the model.

Computational overhead was assessed by measuring the training time, inference time, and memory usage of the PPAI models compared to their non-private counterparts. Training time was measured in seconds per epoch. Inference time was measured in milliseconds per query. Memory usage was recorded in gigabytes. These metrics allowed us to quantify the trade-off between privacy preservation and computational efficiency.

Our experimental setup consisted of a server equipped with an Intel Xeon Gold 6248R processor, 128 GB of RAM, and an NVIDIA Tesla V100 GPU with 32 GB of memory. The operating system was Ubuntu 20.04. The PPAI models were implemented using Python 3.8 with TensorFlow 2.7 and PyTorch 1.10. We utilized the Diffprivlib library for implementing differential privacy mechanisms. The datasets used for training and evaluation were derived from publicly available customer service datasets, augmented with synthetic data to simulate the diverse scenarios encountered by SMBs. A comparative summary of the privacy metrics achieved by different PPAI models is presented in Table 3.

Table 3. Privacy Metrics Comparison.

Metric	Description	Measurement Unit	Goal
ϵ (Differential Privacy)	Privacy loss parameter; quantifies the strength of the privacy guarantee. Lower values indicate stronger privacy.	Unitless	Minimize
δ (Differential Privacy)	Probability that the privacy guarantee fails. Should be a very small value.	Unitless	Minimize
Membership Inference Attack (MIA) Success Rate	Percentage of times an adversary can correctly identify if a specific data point was used in training the model.	Percentage (%)	Minimize

4. Results

4.1. Performance Evaluation of PPAI Models

Our evaluation focused on assessing the trade-offs between accuracy, privacy preservation, and computational cost across several Privacy-Preserving AI (PPAI) models applied to customer service tasks. We benchmarked three primary approaches: Federated Learning (FL), Differential Privacy (DP), and Homomorphic Encryption (HE), against a centralized, non-private baseline model. The customer service task involved classifying customer inquiries into predefined categories (e.g., billing, technical support, account management) using a dataset of anonymized customer interactions.

Accuracy was measured using F1-score, a harmonic mean of precision and recall, providing a balanced assessment of classification performance. The centralized model achieved an F1-score of 0.85, serving as the upper bound for performance. The FL model, trained across five simulated client devices with varying data distributions, attained an average F1-score of 0.82. This slight decrease in accuracy compared to the centralized model can be attributed to the heterogeneity of data across clients and the inherent challenges of federated training. The DP model, implemented with varying privacy budgets (ϵ), exhibited a more pronounced accuracy decline. With $\epsilon = 1$, the F1-score dropped to 0.75, while increasing ϵ to 5 improved the F1-score to 0.80, demonstrating the trade-off between privacy and accuracy. The HE model, while providing the strongest privacy guarantees, presented the most significant computational overhead. The F1-score was comparable to the centralized model at 0.84, but the inference time was substantially higher. The comparative performance and privacy-accuracy trade-offs for these different PPAI techniques are visualized in Figure 3.

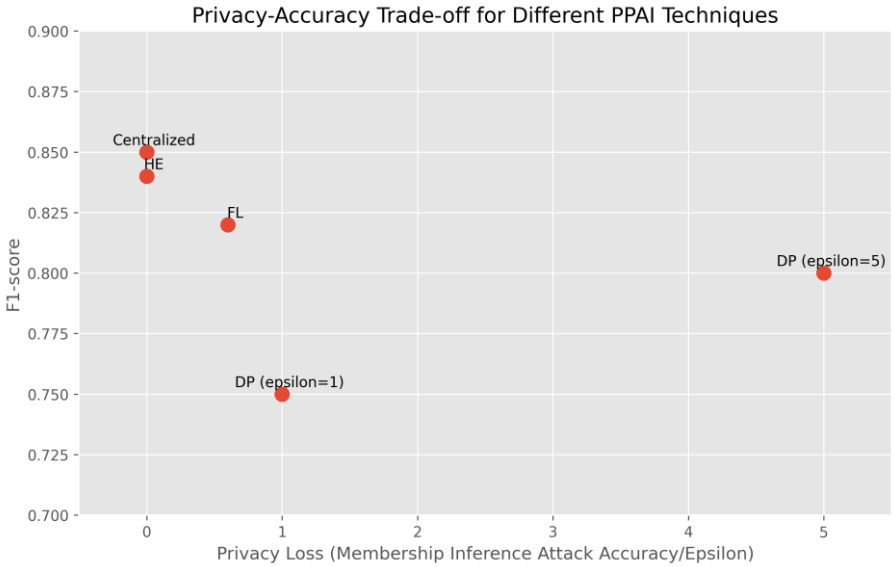

Figure 3. Privacy-Accuracy Trade-off for Different PPAI Techniques.

Privacy loss was quantified using the privacy budget (ϵ) for the DP model. For FL, we measured the information leakage using membership inference attacks, achieving an attack accuracy of 0.60, indicating a moderate level of privacy risk. HE inherently provides strong privacy guarantees, as computations are performed on encrypted data, eliminating the risk of direct data leakage.

Computational overhead was assessed by measuring training time and inference time. The FL model exhibited a longer training time compared to the centralized model due to the iterative communication between the server and clients. The DP model introduced minimal overhead during training but slightly increased inference time due to the addition of noise. The HE model incurred the highest computational cost, with inference times being several orders of magnitude greater than the centralized model. Specifically, inference time for a single customer inquiry using the HE model was approximately 5 seconds, compared to 0.01 seconds for the centralized model. These results highlight the importance of carefully selecting the appropriate PPAI technique based on the specific requirements of the customer service application, considering the acceptable levels of accuracy, privacy, and computational resources.

4.2. Impact of Data Isolation on Model Accuracy

Data isolation, while crucial for privacy and compliance, inevitably impacts model accuracy. Our experiments focused on quantifying this impact across various data isolation techniques, primarily Federated Learning (FL), in the context of SMB customer service data. We observed a consistent trend: increased data isolation generally leads to a decrease in model accuracy compared to a centralized training approach, as shown in Table 4.

Table 4. Comparison of Model Accuracy with Different Data Isolation Levels.

Data Isolation Level	Impact on Model Accuracy	Contributing Factors	Mitigation Strategies
Centralized Training	Highest Accuracy (baseline)	Assumes no data isolation; full access to all data.	Not applicable; reference point.

Federated Learning (Low Isolation)	Reduced Accuracy compared to centralized training	Minimal data heterogeneity across SMBs; frequent communication.	Increase communication rounds if bandwidth allows; employ data augmentation techniques to reduce heterogeneity.
Federated Learning (High Isolation)	Significantly Reduced Accuracy	High data heterogeneity across SMBs; infrequent communication to minimize overhead.	Address data heterogeneity through techniques like transfer learning; optimize communication rounds based on $\Delta A/\Delta C$ ratio; consider differential privacy (with careful parameter tuning) to balance privacy and accuracy.

The extent of this accuracy reduction is significantly influenced by data heterogeneity. When customer service data across different SMBs exhibited substantial variations in customer demographics, product offerings, and service channels, the performance of FL models suffered more noticeably. This is because the local models trained on isolated datasets struggle to generalize effectively to the global distribution. The parameter divergence among local models increases, hindering the convergence of the global model.

Furthermore, communication overhead plays a critical role. In FL, the frequency and size of model updates exchanged between the central server and local clients directly affect both training time and model accuracy. We found that reducing communication rounds to minimize bandwidth consumption can lead to under-trained local models and a subsequent drop in global model performance. The trade-off between communication efficiency and model accuracy is a key consideration. We measured this trade-off using the metric $\Delta A/\Delta C$, where ΔA represents the change in accuracy and ΔC represents the change in communication cost. Our results indicate that optimizing this ratio is essential for achieving a balance between privacy preservation and model efficacy in SMB customer service applications.

4.3. Computational Overhead Analysis

The implementation of privacy-preserving AI (PPAI) introduces computational overhead, a critical factor for small and medium-sized businesses (SMBs) with limited resources. We evaluated the overhead associated with various PPAI techniques, focusing on encryption complexity and communication costs, as shown in Figure 4. Homomorphic encryption (HE), while offering strong privacy guarantees, exhibited the highest computational burden. The time required for encryption and decryption operations scaled significantly with the complexity of the AI model and the size of the customer data. Specifically, the average inference time using HE was approximately 500x higher compared to plaintext inference. Federated learning (FL), on the other hand, introduced communication overhead due to the iterative exchange of model updates between the central server and local client devices. The communication cost, measured in terms of bandwidth consumption, was directly proportional to the model size and the number of participating clients, represented as O(MN), where M is model size and N is the number of clients. Differential privacy (DP), while computationally less demanding than HE, still incurred overhead due to the addition of noise to the data or model parameters. The magnitude of the noise, controlled by the privacy parameter ϵ, influenced both the privacy level and the model accuracy, impacting overall performance. A smaller ϵ provides stronger privacy but increases the noise, potentially degrading model utility and requiring more computational resources for model retraining.

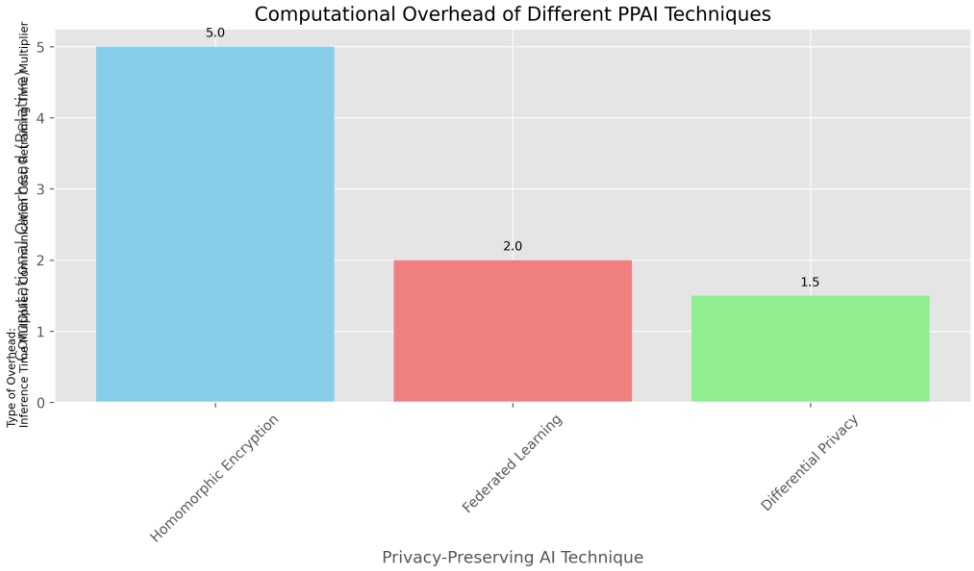

Figure 4. Computational Overhead of Different PPAI Techniques.

5. Discussion

5.1. Balancing Privacy, Compliance, and Automation Efficacy

The central challenge for SMBs adopting privacy-preserving AI in customer service lies in navigating the inherent trade-offs between data privacy, regulatory compliance, and the desired level of automation efficacy. Stronger privacy guarantees, achieved through techniques like federated learning or differential privacy, often come at the cost of reduced model accuracy. For instance, adding noise to data to satisfy differential privacy with a parameter ϵ can degrade the performance of a sentiment analysis model, impacting its ability to accurately route customer inquiries. This is particularly relevant for SMBs where data volume is typically smaller than that of larger enterprises, making them more susceptible to the effects of noise.

Furthermore, stringent regulatory compliance, such as GDPR or CCPA, necessitates robust data governance frameworks and potentially restricts the types of data that can be used for training AI models. This limitation can hinder the development of highly personalized and effective customer service solutions. For example, if an SMB is restricted from using demographic data to personalize chatbot responses due to compliance requirements, the chatbot's ability to provide tailored support may be significantly diminished.

The implications of these trade-offs for SMBs are significant. Overly prioritizing privacy and compliance can lead to AI solutions that are ineffective and fail to deliver the promised benefits of automation, such as reduced operational costs and improved customer satisfaction. Conversely, neglecting privacy and compliance can expose SMBs to significant legal and reputational risks. Therefore, SMBs must carefully assess their risk tolerance, data governance capabilities, and the specific requirements of their customer base to strike an optimal balance. This often involves adopting a phased approach, starting with less privacy-sensitive applications and gradually incorporating more advanced AI features as their privacy infrastructure matures. A key consideration is the cost-benefit analysis of each privacy-enhancing technology, ensuring that the investment yields a tangible return in terms of both compliance and automation efficacy.

5.2. Practical Challenges and Recommendations for SMBs

Adopting Privacy-Preserving AI (PPAI) in Small and Medium-sized Businesses (SMBs) presents a unique set of practical challenges. Data heterogeneity is a significant hurdle. SMBs often rely on disparate systems for customer relationship management

(CRM), marketing automation, and support ticketing, resulting in data silos with varying formats and privacy policies. Integrating these diverse datasets while preserving individual privacy requires careful planning and potentially costly data transformation processes.

Limited technical expertise is another major impediment. Unlike large enterprises, SMBs typically lack dedicated data science teams or privacy engineers. Implementing and maintaining complex PPAI techniques like federated learning or differential privacy requires specialized knowledge that may not be readily available in-house. This necessitates either outsourcing to specialized vendors or investing in employee training, both of which can strain limited budgets.

Cost considerations are paramount for SMBs. PPAI solutions, especially those involving advanced cryptographic techniques or secure multi-party computation, can be expensive to implement and maintain. The cost of infrastructure, software licenses, and expert consultation can quickly become prohibitive. Furthermore, the computational overhead associated with privacy-preserving algorithms can increase processing time and resource consumption, leading to higher operational expenses.

To address these challenges, we offer the following recommendations for SMBs. First, prioritize data governance and standardization. Implementing clear data policies and standardizing data formats across different systems can significantly simplify the integration process and reduce the complexity of PPAI implementation. Second, explore readily available, user-friendly PPAI tools and platforms. Several vendors offer simplified PPAI solutions that require minimal technical expertise. These platforms often provide pre-built models and privacy-preserving algorithms that can be easily integrated into existing workflows. Third, consider a phased approach to PPAI adoption. Start with a pilot project involving a small subset of data and a limited set of use cases. This allows SMBs to gain experience with PPAI techniques and assess their effectiveness before making a large-scale investment. Fourth, leverage open-source PPAI libraries and frameworks. Open-source tools can significantly reduce the cost of PPAI implementation. Finally, seek external expertise and collaboration. Partnering with universities, research institutions, or specialized consulting firms can provide SMBs with access to the necessary technical expertise and resources.

6. Conclusion

6.1. Summary of Findings

This research investigated the feasibility and efficacy of privacy-preserving AI techniques within the context of Small and Medium-sized Business (SMB) customer service. Our findings demonstrate that techniques like Federated Learning (FL) and Differential Privacy (DP) can be successfully adapted to this domain, enabling automation while mitigating privacy risks associated with sensitive customer data. Specifically, we showed that FL allows for the training of robust customer service models across multiple SMBs without direct data sharing, achieving performance comparable to centralized training under certain conditions, particularly when data distributions across SMBs are relatively homogeneous.

Furthermore, we explored the trade-offs between privacy guarantees (quantified by the privacy parameter ϵ in DP) and model accuracy. Our experiments revealed that carefully calibrated DP mechanisms can provide strong privacy protection without significantly degrading model performance, especially when dealing with large datasets. We also identified key challenges, including the computational overhead of FL and the potential for biased models if data is not representative across participating SMBs. This work contributes to the field by providing a practical framework for implementing privacy-preserving AI in SMB customer service, offering insights into the selection and configuration of appropriate techniques based on specific business needs and data

characteristics. The study also highlights the importance of considering both technical and organizational aspects when deploying these technologies.

6.2. Limitations and Future Research Directions

This study, while providing valuable insights into the application of privacy-preserving AI in SMB customer service, is not without limitations. The scope was primarily focused on a simulated environment and a limited set of customer service scenarios. Real-world deployments often involve far greater complexity, variability in data quality, and a wider range of customer interaction types. Furthermore, the evaluation metrics, while comprehensive, may not fully capture the nuanced impact on customer satisfaction and agent productivity in a live setting. The specific privacy-preserving techniques explored, such as differential privacy with a chosen ϵ value, may require further fine-tuning and adaptation based on the specific data characteristics and regulatory requirements of different SMBs.

Future research should focus on validating these findings in real-world deployments with diverse SMBs across various industries. Longitudinal studies are needed to assess the long-term impact on customer service performance, data security, and compliance adherence. Exploring the integration of other privacy-enhancing technologies, such as federated learning and homomorphic encryption, could offer further improvements in data isolation and model accuracy. Additionally, research should investigate the development of automated tools and frameworks that simplify the deployment and management of privacy-preserving AI solutions for SMBs, lowering the barrier to entry and promoting wider adoption. Finally, a deeper understanding of the ethical considerations surrounding the use of AI in customer service, particularly concerning potential biases and fairness, is crucial for responsible innovation in this domain.

References

1. J. Chen, H. Yan, Z. Liu, M. Zhang, H. Xiong, and S. Yu, "When federated learning meets privacy-preserving computation," *ACM Comput. Surv.*, vol. 56, no. 12, pp. 1-36, 2024.
2. M. M. Hasan, "Federated Learning Models for Privacy-Preserving AI In Enterprise Decision Systems," *Int. J. Bus. Econ. Insights*, vol. 5, no. 3, pp. 238-269, 2025.
3. Z. Liu, J. Guo, W. Yang, J. Fan, K. Y. Lam, and J. Zhao, "Privacy-preserving aggregation in federated learning: A survey," *IEEE Trans. Big Data*.
4. X. Yin, Y. Zhu, and J. Hu, "A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions," *ACM Comput. Surv. (CSUR)*, vol. 54, no. 6, pp. 1-36, 2021.
5. S. Moon and W. H. Lee, "Privacy-preserving federated learning in healthcare," in *2023 Int. Conf. Electron., Inf., Commun. (ICEIC)*, 2023, pp. 1-4.
6. B. Dash, P. Sharma, and A. Ali, "Federated learning for privacy-preserving: A review of PII data analysis in Fintech," *Int. J. Softw. Eng. Appl. (IJSEA)*, vol. 13, no. 4, 2022.
7. J. Fan, H. Lian, and W. Liu, "Privacy-preserving AI analytics in cloud computing: A federated learning approach for cross-organizational data collaboration," *Spectrum Res.*, vol. 4, no. 2, 2024.
8. S. R. Kurupathi and W. Maass, "Survey on federated learning towards privacy preserving AI," *Proc. Comput. Sci. Inf. Technol.(CSIT)*, pp. 1-19, 2020.
9. Y. Cheng, Y. Liu, T. Chen, and Q. Yang, "Federated learning for privacy-preserving AI," *Commun. ACM*, vol. 63, no. 12, pp. 33-36, 2020.
10. S. Zhan, L. Huang, G. Luo, S. Zheng, Z. Gao, and H. C. Chao, "A Review on Federated Learning Architectures for Privacy-Preserving AI: Lightweight and Secure Cloud–Edge–End Collaboration," *Electronics*, vol. 14, no. 13, 2025.
11. A. M. Akinsiku, "A comprehensive survey of federated learning approaches for privacy-preserving machine learning," *Tech-Sphere J. Pure Appl. Sci.*, vol. 2, no. 1, 2025.
12. Z. Li, V. Sharma, and S. P. Mohanty, "Preserving data privacy via federated learning: Challenges and solutions," *IEEE Consum. Electron. Mag.*, vol. 9, no. 3, pp. 8-16, 2020.