Journal of Applied Statistics & Machine Learning
Vol. 3, Nos. 1-2, 2024, pp. 85-101, ISSN: 2583-2891
© Black Rose Publication
URL : www.blackroseindia.com/jasml

Privacy-Preserving Customer Support: A Framework for Secure and Scalable Interactions

Anant P. Awasthi[1], Girdhar G. Agarwal[5], Chandraketu Singh[2], Rakshit Varma[3], Sanchit Sharma[4]

[1, 3, 4]Optum Global Solutions, Noida, Gautam Buddha Nagar, Uttar Pradesh, India
[2]Jaipuria Institute of Management, Lucknow, Uttar Pradesh, India
[5]Ex-Professor, University of Lucknow, Lucknow, Uttar Pradesh, India
Corresponding author E-mail: anant.awasthi@outlook.com

***Abstract**:* Artificial intelligence (AI) has revolutionized customer support by improving efficiency and enhancing user experiences. However, traditional machine learning (ML) methods often rely on extensive local training with sensitive data, raising significant privacy concerns and posing challenges for compliance with regulations like the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). While privacy-preserving methods such as anonymization, differential privacy, and federated learning address some of these concerns, they often struggle with scalability, utility, and implementation complexity.

This paper presents the Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework—a groundbreaking approach that harnesses the power of large language models (LLMs) in zero-shot learning. Unlike traditional ML techniques, PP-ZSL avoids local data training by leveraging pre-trained LLMs to generate accurate responses directly. The framework integrates real-time data anonymization to protect sensitive information, retrieval-augmented generation (RAG) for handling domain-specific queries, and advanced post-processing mechanisms to ensure compliance with privacy regulations. Together, these elements reduce risks, simplify compliance, and enhance scalability and efficiency.

Our analysis reveals that the PP-ZSL framework delivers privacy-compliant, accurate responses while minimizing the cost and complexity of deploying AI-driven customer support systems. The framework holds transformative potential for industries such as financial services, healthcare, e-commerce, legal support, telecommunications, and government services. By balancing privacy with performance, PP-ZSL sets the stage for secure, efficient, and regulation-ready AI solutions in customer interactions.

Keywords: Privacy-Preserving Zero-Shot Learning (PP-ZSL), Large Language Models (LLMs), Data Privacy and Anonymization, Zero-Shot Learning (ZSL), Regulatory Compliance (GDPR, CCPA), Retrieval-Augmented Generation (RAG)

TO CITE THIS ARTICLE

Awasthi, A.P., Agarwal, G.G., Singh, C., Varma, R., & Sharma, S. (2024). Privacy-Preserving Customer Support: A Framework for Secure and Scalable Interactions, *Journal of Applied Statistics & Machine Learning,* 3(1-2), pp. 85-101.

1. Introduction

Artificial intelligence (AI) has revolutionized customer support, enabling real-time, intelligent assistance that meets evolving customer needs. However, these advancements come with a critical challenge: safeguarding data privacy. Traditional machine learning (ML) models often depend on extensive local training using organization-specific datasets, which frequently include sensitive information such as personally identifiable information (PII), financial records, or contractual details. This reliance heightens privacy risks and complicates compliance with stringent regulations like the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

The local training approach introduces significant vulnerabilities in data handling, storage, and sharing, increasing the risk of data breaches and misuse. Existing privacy-preserving methods, such as anonymization, differential privacy, and federated learning, provide partial solutions but come with notable drawbacks. Anonymization can diminish the usability of data, differential privacy often results in trade-offs that affect model performance, and federated learning poses operational challenges, including scalability and maintenance issues. As customers demand greater privacy and seamless service, organizations face a pressing need for innovative solutions that address these challenges without sacrificing service quality.

This paper introduces a transformative approach by leveraging Zero-Shot Learning (ZSL) through large language models (LLMs). Unlike conventional ML methods, which require fine-tuning or local training, ZSL utilizes pre-trained LLMs to generate responses directly, eliminating the need for sensitive data to be stored or processed locally. This inherent reduction in data handling risk makes ZSL an attractive, scalable, and cost-effective solution for delivering secure and accurate customer support.

To operationalize this concept, we propose the Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework, which integrates several privacy-focused components:

1. **Real-Time Data Anonymization**: Sensitive details, such as PII and financial information, are redacted or masked before the query reaches the LLM.

2. **Retrieval-Augmented Generation (RAG)**: For domain-specific queries, the framework retrieves information from secure, non-sensitive knowledge repositories.

3. **Post-Processing and Validation**: Generated responses are rigorously validated to ensure compliance with privacy policies and regulatory standards.

The PP-ZSL framework not only reduces the risks inherent in traditional ML models but also simplifies regulatory compliance by eliminating the need for local data storage and training. This approach streamlines processes like audits and ensures adherence to key privacy principles, including the "right to be forgotten" and data minimization. Furthermore, the framework is designed to meet operational demands, addressing challenges such as latency, accuracy, and scalability in customer support systems.

In this paper, we examine the limitations of existing privacy-preserving methods, detail the design and implementation of the PP-ZSL framework, and evaluate its performance through empirical case studies. Our findings demonstrate that PP-ZSL enables organizations to achieve a crucial balance between cutting-edge AI performance and robust privacy protections. In doing so, it lays the groundwork for ethical, secure, and efficient AI-powered customer support solutions.

2. Review of Literature

The intersection of machine learning, privacy preservation, and customer support has been a focal point of research across diverse fields. This review synthesizes key contributions in privacy-preserving techniques, the evolution of large language models (LLMs), and their application in customer support, identifying critical gaps that this publication aims to address.

2.1. Privacy Challenges in Machine Learning

Machine learning models thrive on large volumes of data for training, but this dependence brings significant privacy concerns, particularly when dealing with sensitive information such as personally identifiable information (PII) or financial records. Traditional anonymization techniques, like pseudonymization, have been shown to offer limited protection due to the risk of re-identification (Rocher et al., 2019). Differential privacy, introduced by Abadi et al. (2016), provides a robust theoretical foundation for preventing data reconstruction but often requires a trade-off with model utility, reducing performance in practical applications. Federated learning (Kairouz et al., 2021) addresses some of these concerns by decentralizing training to local devices, yet it comes with its own challenges, including communication overhead and scalability issues. These limitations highlight the need for alternative methods that prioritize privacy without compromising on model performance or scalability.

2.2. Regulatory Landscape and Privacy Compliance

Privacy regulations, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), have significantly influenced the way businesses handle data. These frameworks emphasize principles like data minimization, the right to be forgotten, and accountability in automated decision-making (European Union, 2016; CCPA, 2018). While they establish a legal mandate for implementing privacy-preserving practices, businesses—especially those in customer support—often struggle to maintain compliance while delivering high-quality service. Existing solutions, though well-intentioned, frequently fall short in balancing regulatory demands with the need for operational efficiency and user satisfaction.

2.3. Large Language Models and Zero-Shot Learning

Recent advancements in large language models (LLMs), such as GPT-3 and T5, have redefined the boundaries of AI by demonstrating the ability to perform complex tasks without task-specific training—a phenomenon known as zero-shot learning (Brown et al., 2020; Raffel et al., 2020). These models, pre-trained on vast and diverse datasets, exhibit impressive generalization capabilities across a wide range of domains. Zero-shot learning presents a promising pathway for privacy preservation by reducing the dependency on local training data, thereby minimizing risks associated with sensitive data handling. However, applying these models to domain-specific scenarios without additional fine-tuning remains a significant challenge, as observed by Zhang et al. (2020). This gap calls for innovative strategies to adapt LLMs to specific applications while maintaining their privacy-preserving advantages.

2.4. Privacy-Preserving LLM Architectures

Recent studies have explored ways to integrate LLMs with privacy-preserving techniques, leading to encouraging developments. For example, retrieval-augmented generation (RAG) allows models to fetch domain-specific knowledge from secure, external repositories, bypassing the need for sensitive data training (Lewis et al., 2020). Additionally, mechanisms like Named Entity Recognition (NER) are employed to anonymize input queries, safeguarding PII before processing (Li et al., 2021). Despite these advancements, a comprehensive framework that seamlessly combines privacy preservation, regulatory compliance, and operational scalability for customer support applications remains unexplored. Addressing this gap is critical for enabling ethical, secure, and efficient AI-driven customer interactions.

2.5. Gaps and Contributions

Despite notable advancements in privacy-preserving methods and the growing adoption of zero-shot learning with large language models (LLMs), there remains a lack of cohesive frameworks that integrate these capabilities for customer support. Existing research tends to focus on isolated aspects, such as anonymization, retrieval-augmented techniques, or regulatory compliance, without addressing how these elements can work together to create a comprehensive solution.

This framework bridges these gaps by introducing a unified framework that combines zero-shot learning, real-time data anonymization, and retrieval-augmented generation (RAG) techniques. The proposed Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework ensures privacy compliance while maintaining scalability, efficiency, and high service quality in customer interactions. By integrating these components, the framework not only addresses existing challenges but also lays the groundwork for secure, ethical, and regulation-ready AI-driven customer support systems.

3. Privacy Challenges in Traditional Approaches

Traditional machine learning (ML) systems depend heavily on vast amounts of training data to achieve high accuracy and task-specific performance. In the context of customer support, this often requires processing sensitive information, such as personally identifiable information (PII), financial records, and contractual agreements. While this data is crucial for building and fine-tuning effective models, it also introduces significant privacy risks. Improper handling of such data can lead to breaches, unauthorized access, and non-compliance with stringent regulations like the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) (European Union, 2016; CCPA, 2018). These challenges place organizations under growing pressure to adopt privacy-preserving strategies without compromising the quality of their ML applications.

One widely used privacy-preserving method is data anonymization, which involves removing or masking identifiable elements from datasets. While anonymization can reduce the risk of exposure, research indicates that anonymized data is often vulnerable to re-identification through auxiliary datasets or sophisticated statistical techniques (Rocher et al., 2019). This limitation compromises its effectiveness, particularly when dealing with highly sensitive or complex datasets. Furthermore, anonymization often diminishes the utility of the data, resulting in reduced model performance. This trade-off makes it less suitable for demanding applications like customer support, where accuracy and reliability are paramount.

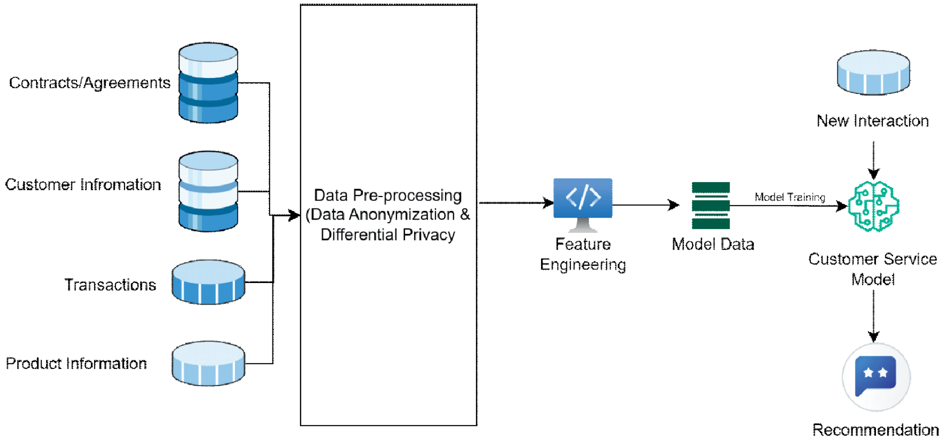

Fig. 3.1: Traditional Machine Learning Solution Development Approach

Differential privacy offers a robust theoretical framework for safeguarding individual contributions within a dataset by introducing controlled noise to data or model outputs (Abadi et al., 2016). This ensures that individual records remain indistinguishable, even when analyzed by sophisticated methods. However, the practical implementation of differential privacy faces a significant challenge: balancing privacy with utility. High levels of noise, while enhancing privacy, can severely degrade model performance, resulting in inaccurate or unhelpful outputs—an unacceptable outcome for customer support applications. Additionally, differential privacy often demands substantial computational resources, making it costly and complex to deploy at scale.

Federated learning has emerged as another promising approach, enabling models to train locally on user devices without transferring sensitive data to a central server (Kairouz et al., 2021). While this decentralized method reduces risks associated with data transmission, it introduces its own set of challenges. Synchronizing models across devices creates communication overhead, and the reliance on local devices makes the system vulnerable to security threats. Moreover, federated learning does not inherently prevent sensitive information from being encoded into the model, leaving the door open to potential privacy breaches.

These limitations highlight the need for innovative solutions that not only reduce reliance on sensitive data but also ensure compliance with privacy regulations. Traditional approaches, though effective in specific contexts, are often resource-intensive, compromise data utility, or fall short of providing comprehensive privacy guarantees. To address these challenges, this paper proposes a zero-shot learning (ZSL) framework that leverages large language models (LLMs) to generate accurate and privacy-compliant outputs without requiring local training on sensitive data. This approach significantly reduces privacy risks and simplifies regulatory compliance.

4. Zero-Shot Learning (ZSL) as a Paradigm Shift

Zero-Shot Learning (ZSL) represents a groundbreaking shift in natural language processing (NLP), enabling AI models to perform tasks without the need for task-specific training. This capability is particularly valuable in scenarios where collecting and annotating data is impractical, expensive, or poses significant privacy concerns. Advances in large language models (LLMs), such as GPT-3 and T5, have demonstrated ZSL's potential to generalize across a wide range of domains using the knowledge acquired during extensive pre-training on diverse datasets (Brown et al., 2020; Raffel et al., 2020). For customer support, ZSL offers a significant leap forward by eliminating the need to train models with organization-specific, and often sensitive, data, thereby enhancing both operational efficiency and data privacy.

Unlike traditional machine learning (ML) approaches, which rely on fine-tuning with domain-specific datasets, ZSL models harness their pre-trained generalization capabilities to handle a wide array of tasks. This makes ZSL particularly well-suited for customer support, where queries often span multiple domains and require dynamic, adaptable responses. ZSL models accomplish this by interpreting task instructions and contextualizing them within their pre-existing knowledge base, often formatted as natural language prompts. For example, a model can respond to a question about financial policies without ever accessing or being trained on an organization's proprietary data. This approach significantly mitigates risks associated with data breaches or unintended exposure (Brown et al., 2020).

By leveraging ZSL, organizations can maintain high levels of accuracy and relevance in customer interactions while minimizing data handling risks and ensuring compliance with privacy regulations. This paradigm shift underscores the potential of LLM-driven ZSL frameworks to redefine how AI systems operate in sensitive and high-stakes environments like customer support.

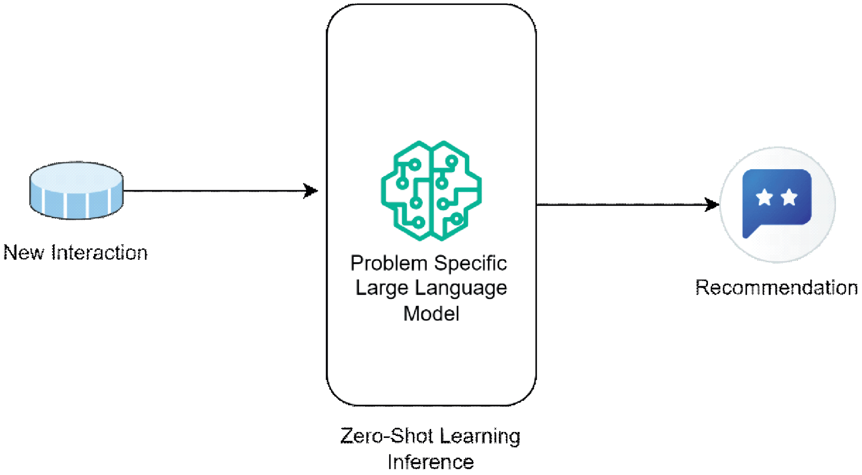
New Interaction
Problem Specific Large Language Model
Zero-Shot Learning Inference
Recommendation

Fig. 4.1: Zero Shot learning (Inference Process) Representation

An additional strength of Zero-Shot Learning (ZSL) is its ability to simplify compliance with privacy regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Traditional machine learning (ML) models often require the storage and processing of sensitive data locally, making it difficult to meet regulatory requirements like the "right to be forgotten" and data minimization principles. In contrast, ZSL-based systems operate without needing sensitive data for training or storage, inherently aligning with these privacy mandates (European Union, 2016; CCPA, 2018).

Despite its advantages, ZSL is not without its challenges, particularly in domain-specific applications. While pre-trained models excel at general tasks, their performance may decline when addressing highly specialized queries or niche topics. To overcome these limitations, hybrid approaches like Retrieval-Augmented Generation (RAG) have been introduced. RAG enhances ZSL capabilities by incorporating external, domain-specific knowledge bases, enabling accurate responses without requiring additional training on sensitive data (Lewis et al., 2020).

The integration of ZSL in customer support represents a paradigm shift, transforming not only operational workflows but also the standard for privacy preservation in AI-driven systems. By eliminating the need for local training and embedding privacy-aware mechanisms, such as dynamic data redaction and external knowledge retrieval, ZSL emerges as a cornerstone for scalable, secure, and compliant AI solutions. This approach not only addresses regulatory complexities but also sets a new benchmark for ethical and efficient AI applications in customer interactions.

5. Proposed Framework: Privacy-Preserving Zero-Shot Learning (PP-ZSL)

The Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework is designed to provide accurate, scalable, and privacy-compliant responses in customer support. It ensures sensitive information is protected throughout the interaction pipeline. Below is a detailed discussion of each component of the framework, along with the flowchart provided below.

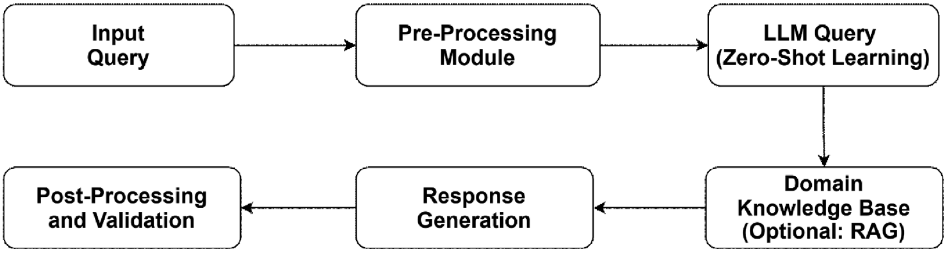
Input Query
Pre-Processing Module
LLM Query (Zero-Shot Learning)
Domain Knowledge Base (Optional: RAG)
Response Generation
Post-Processing and Validation

Fig. 5.1: Proposed Privacy-Preserving Zero-shot Learning Framework

5.1. Input Query

The framework starts when a customer submits a query. This input often contains sensitive information, such as personally identifiable information (PII), financial data, or contractual details, which must be protected from the very beginning. The framework ensures that this data remains secure throughout the interaction pipeline, preventing unauthorized access. This initial step sets the stage for privacy-aware processing in compliance with regulations like GDPR and CCPA.

5.2. Pre-Processing Module

The pre-processing module is the first layer of privacy protection. It identifies and anonymizes sensitive data in the input query before it reaches the language model.

Techniques like Named Entity Recognition (NER) are used to detect sensitive elements such as names, account numbers, and dates. These elements are either masked or replaced with placeholders using tokenization and redaction. The module also includes dynamic anonymization, which adjusts the level of redaction based on specific privacy policies or organizational needs. By processing the query at this stage, the framework ensures that raw sensitive data is not passed to the LLM, reducing the risk of inadvertent data exposure.

5.3. LLM Query (Zero-Shot Learning)

At the core of the framework is the use of pre-trained Large Language Models (LLMs) in zero-shot learning (ZSL) mode. Unlike traditional machine learning methods, which require fine-tuning with organization-specific data, ZSL relies on the generalization capabilities of pre-trained models to interpret and respond to queries. This eliminates the need for local training, significantly reducing privacy risks and operational costs. The LLM uses its extensive pre-trained knowledge to handle a wide range of customer queries, ensuring accurate and contextually relevant responses without needing sensitive data for additional training.

5.4. Domain Knowledge Base (Optional: RAG)

To improve domain-specific accuracy, the framework can optionally use a retrieval-augmented generation (RAG) module. This module supplements the LLM’s zero-shot capabilities by retrieving relevant, non-sensitive information from secure knowledge repositories. For example, queries about company policies or technical details can be addressed by fetching up-to-date information from external databases or APIs. This ensures the LLM delivers precise and contextually accurate responses without requiring sensitive data storage or processing locally.

5.5. Response Generation

Once the query is processed, the LLM generates a response based on the anonymized input and any additional information retrieved through the RAG module. The generated response is designed to address the customer’s query while preserving the anonymization applied during pre-processing. This ensures that sensitive data is not reintroduced into the response, maintaining privacy throughout the interaction.

5.6. Post-Processing and Validation

The final step involves post-processing and validation to ensure the generated response complies with privacy regulations and organizational standards. Privacy filters are applied to detect and remove any accidentally reintroduced sensitive data. A compliance audit is conducted to confirm the response meets requirements such as GDPR's "right to be forgotten" and CCPA's data minimization principles. This guarantees that the final response is privacy-compliant and ready for delivery to the customer.

The PP-ZSL framework brings together these components into a seamless workflow, effectively addressing the dual challenges of data privacy and operational scalability. By anonymizing queries, utilizing zero-shot learning, and applying thorough validation, the framework provides a transformative solution for AI-driven customer support that ensures accuracy, efficiency, and compliance.

The framework guarantees end-to-end privacy preservation while delivering accurate and efficient responses. It starts with query anonymization, processes the input using pre-trained LLMs with optional domain-specific augmentation, and concludes with post-processing validation to ensure compliance with privacy regulations.

This modular design allows organizations to implement the framework flexibly and adapt it to meet specific operational or regulatory requirements.

6. Future Directions

The proposed Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework offers a solid foundation for tackling privacy challenges in customer support. However, there are several areas for future research and development that could improve its scalability, adaptability, and overall effectiveness. This section highlights key directions, including hybrid approaches, real-time privacy enhancements, contextual understanding, and compliance with emerging global privacy laws.

6.1. Hybrid Approaches for Domain-Specific Adaptability

While zero-shot learning (ZSL) is effective for general tasks, its performance may decline in highly specialized contexts. Future research could focus on hybrid models that combine ZSL with lightweight fine-tuning on anonymized or synthetic data. Methods like prompt engineering (Brown et al., 2020) or instruction-tuning (Wei et al., 2022) could help bridge the gap between generalization and domain-specific accuracy. Synthetic datasets created using privacy-preserving techniques, such as differential privacy (Abadi et al., 2016) or generative adversarial networks (GANs) (Goodfellow et al., 2014), could enable safe task-specific training without exposing sensitive information.

6.2. Advances in Real-Time Privacy Filters

Enhancing real-time data anonymization and redaction tools is essential for practical deployment. Named Entity Recognition (NER) models could be improved to detect and anonymize sensitive information with greater accuracy and contextual awareness (Li et al., 2021). Future efforts could also explore multi-modal privacy filters capable of processing textual, visual, and auditory data to support diverse customer support channels. Using on-device processing for privacy filtering could further minimize latency and improve data security.

6.3. Enhancing Contextual Understanding in LLMs

Current large language models (LLMs) often struggle with maintaining context in multi-turn conversations. Future research could incorporate context-tracking mechanisms to allow models to securely retain and use dialogue history. Techniques like retrieval-augmented generation (RAG) (Lewis et al., 2020) could dynamically fetch relevant contextual information, reducing dependence on sensitive local datasets.

6.4. Addressing Scalability and Cost-Efficiency

Scaling ZSL frameworks for real-time customer support systems presents challenges due to computational demands. Future research could explore optimization techniques like quantization (Gong et al., 2014) and pruning (Han et al., 2015) to lower inference costs without compromising performance. Developing serverless architectures that offload model execution to edge devices could also balance scalability and cost.

6.5. Adapting to Emerging Privacy Regulations

As global privacy laws evolve, ensuring that the PP-ZSL framework remains compliant with new regulations is critical. Future research could focus on creating dynamic compliance modules that adjust to regional privacy requirements in real time. For instance, adapting to laws like India's Data Protection Act or China's Personal Information

Protection Law requires localized solutions that balance compliance and operational efficiency (Chik, 2021). Automated compliance monitoring systems could flag potential violations in model outputs, ensuring continued adherence to legal standards.

6.6. Ethical Considerations and Bias Mitigation

Biases inherited from pre-training datasets can lead to discriminatory or unethical outputs (Binns et al., 2018). Future work could explore methods for detecting and mitigating bias in LLMs while adhering to privacy-preserving principles. Techniques like adversarial training (Zhang et al., 2018) or counterfactual data augmentation (Kaushik et al., 2020) could improve fairness without compromising sensitive information.

6.7. Explainability and Trust in AI

Building trust in AI systems requires transparency and explainability. Future research could focus on developing interpretable mechanisms tailored to privacy-preserving frameworks. Techniques such as SHAP (Lundberg & Lee, 2017) or LIME (Ribeiro et al., 2016) could be adapted to explain LLM decisions in real-time customer support interactions while maintaining privacy.

6.8. Exploring Cross-Domain Applications

Although this framework focuses on customer support, its principles can be extended to other areas like healthcare, finance, and legal services. Each domain brings unique privacy challenges and regulatory requirements, providing opportunities for cross-domain adaptation. For example, applying ZSL in healthcare could anonymize patient data while delivering accurate diagnosis and treatment recommendations (Rieke et al., 2020).

7. Conclusion

The Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework represents a transformative approach for AI-ready organizations, providing a scalable, efficient, and privacy-compliant solution for customer support. By eliminating the need for local training on sensitive datasets, the framework drastically reduces data privacy risks while ensuring compliance with regulations such as GDPR and CCPA. Its real-time anonymization safeguards sensitive information, and the zero-shot learning capabilities of pre-trained LLMs enable accurate and contextual responses without requiring additional fine-tuning. This approach not only improves customer satisfaction but also reduces operational complexity, allowing for rapid deployment and cost savings.

With its dynamic knowledge integration through retrieval-augmented generation (RAG), the framework adapts seamlessly to domain-specific requirements, making it suitable for organizations across a wide range of industries. It simplifies compliance reporting by adhering to privacy-by-design principles, lowering legal risks, and preparing businesses for changes in regulatory environments. By emphasizing modularity and innovation, the PP-ZSL framework enables organizations to focus resources on strategic goals while maintaining consistent and high-quality customer support.

In summary, the PP-ZSL framework equips organizations to meet the dual demands of privacy and performance. Whether addressing sensitive financial inquiries, supporting healthcare operations, or managing e-commerce assistance, the framework strikes a balance between operational efficiency, customer experience, and data security. It positions businesses to succeed in an increasingly privacy-conscious digital world.

8. Potential Use Cases

The Privacy-Preserving Zero-Shot Learning (PP-ZSL) framework is highly versatile, with applications spanning industries where customer support demands accuracy, scalability, and strict adherence to privacy regulations.

In **financial services**, the framework can securely address banking and insurance inquiries, such as account details, loan applications, and fraud detection. It ensures that sensitive financial data and personal identifiers are protected while providing real-time, context-aware responses.

In **healthcare**, the framework can assist with general medical advice, appointment scheduling, and clarifications on test results without accessing sensitive health records. For telemedicine, it supports virtual consultations while maintaining patient privacy and confidentiality.

In **e-commerce**, the PP-ZSL framework can handle queries about orders, returns, and personalized product recommendations using anonymized customer preferences, ensuring a secure and seamless shopping experience.

For **legal and contractual support**, the framework can analyze contracts, explain legal clauses, and resolve disputes without exposing sensitive legal or contractual data.

In **telecommunications**, the framework can assist customers with technical issues, billing inquiries, and service outages, all while safeguarding device identifiers and account details.

In **retail and loyalty programs**, it can manage loyalty rewards, promotional offers, and customer feedback securely, protecting identifiable information.

In **government and public services**, the framework supports citizen inquiries about services, benefits, or taxes while ensuring confidentiality of personal and financial data. It can also assist with regulatory compliance, helping citizens navigate policy changes securely.

In **education and e-learning**, it addresses student inquiries about enrollment, courses, and grades and provides learning assistance without compromising personal data.

In **human resources**, the framework supports employee inquiries about payroll, benefits, and leave policies while ensuring privacy. It can also assist in recruitment by providing anonymized feedback on applications and interviews.

In the **travel and hospitality sector**, the framework can handle travel bookings, cancellations, and itinerary changes while protecting sensitive payment and travel details. It also enables secure, personalized travel recommendations.

For **real estate**, the framework facilitates secure property-related queries, pricing information, and mortgage assistance, protecting sensitive financial and personal data.

In **customer feedback and analytics**, it securely aggregates, analyzes, and measures sentiment in customer feedback while ensuring data anonymity.

By complying with privacy laws like GDPR and CCPA, the PP-ZSL framework provides a secure, scalable, and efficient solution for customer support across industries. It minimizes risks associated with handling sensitive data, enhances the customer experience through secure, real-time interactions, and serves as a versatile tool for privacy-focused support in any domain.

9. Conflict of Interest

Authors do not have any conflict of interest as there is no external/internal funding used to complete this work.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.

Binns, R., Veale, M., Van Kleek, M., & Shadbolt, N. (2018). “It’s reducing a human being to a percentage”: Perceptions of justice in algorithmic decisions. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–14.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems, 33*, 1877–1901.

European Union. (2016). *General Data Protection Regulation (GDPR)*. Official Journal of the European Union.

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning, 14*(1–2), 1–210.

California Consumer Privacy Act of 2018, Cal. Civ. Code § 1798.100 (2018).

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., ... & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research, 21*(140), 1–67.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*, 5998–6008.

Woodruff, A., Fox, S. E., Rousso-Schindler, S., & Warshaw, J. (2018). A qualitative exploration of perceptions of algorithmic fairness. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–14.

Zhang, Y., Yu, M., Chang, S., Cheng, Y., & Glass, M. (2020). Query-driven contextual reasoning for zero-shot learning in knowledge bases. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 4413–4424.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Riedel, S. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems, 33*, 9459–9474.

Li, F., Zhang, W., Yu, Y., & Qian, H. (2021). Anonymizing sensitive information in natural language processing: A systematic review. *IEEE Access, 9*, 124623–124639.

Rocher, L., Hendrickx, J. M., & de Montjoye, Y. A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications, 10*(1), 1–9.

Chik, W. B. (2021). Global privacy law trends: GDPR’s influence and the Asia-Pacific region. *Computer Law & Security Review, 40*, 105527.

Gong, Y., Liu, L., Yang, M., & Bourdev, L. (2014). Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*.

Han, S., Pool, J., Tran, J., & Dally, W. J. (2015). Learning both weights and connections for efficient neural networks. *Advances in Neural Information Processing Systems, 28*, 1135–1143.

Kaushik, D., Hovy, E., & Lipton, Z. C. (2020). Learning the difference that makes a difference with counterfactually augmented data. *International Conference on Learning Representations (ICLR)*.

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*, 4765–4774.

Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine, 3*(1), 1–7.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., ... & Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.

Zhang, B. H., Lemoine, B., & Mitchell, M. (2018). Mitigating unwanted biases with adversarial learning. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 335–340.

Jiang, Y., Shao, Y., Ma, D., Semnani, S. J., & Lam, M. S. (2024). Into the unknown unknowns: Engaged human learning through participation in language model agent conversations. *arXiv*. https://arxiv.org/abs/2408.15232

Shao, Y., Jiang, Y., Kanell, T. A., Xu, P., Khattab, O., & Lam, M. S. (2024). Assisting in writing Wikipedia-like articles from scratch with large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.