

# Balancing Privacy and Personalization: AI Solutions for Hyper - Personalized Media Platforms

Raghu K Para

Independent Researcher, Artificial Intelligence, Ontario, Canada

**Abstract:** *Hyper - personalized online media platforms are increasingly leveraging deep learning and natural language processing (NLP) to tailor content, recommendations, and interactions to individual users. However, such personalization often comes at the expense of user privacy due to extensive data collection and centralized model training processes. This study investigates the intersection of AI - driven hyper - personalization and privacy - preserving technologies in online media platforms. It focuses on federated learning and encryption - based NLP systems to address privacy concerns while maintaining personalization efficacy. By exploring cutting - edge privacy - enhancing methods and cryptographic protocols, the paper proposes frameworks to balance these competing objectives. It also examines challenges such as system heterogeneity, computational overhead, and regulatory compliance, offering future directions for secure, scalable, and user - centric AI solutions. Finally, we present future directions for secure collaborative learning, advanced cryptographic approaches, and policy considerations that can help shape an era of user - centric, privacy - preserving AI.*

**Keywords:** privacy - preserving AI, federated learning, hyper - personalization, encryption - based NLP, user data protection

## 1. Introduction

With the proliferation of online media platforms, personalization has become a cornerstone of user engagement. Algorithms that curate articles, videos, advertisements, and social media feeds frequently rely on detailed user profiles gleaned from browsing behavior, content interactions, demographic attributes, and often private data (Kumar et al., 2020). Hyper - personalization raises concerns about data security, as large volumes of user information typically need to be aggregated and centralized for effective machine learning and recommendation systems (Krämer & Böhrs, 2017). In an era marked by stringent data protection regulations, such as the European Union's General Data Protection Regulation

(GDPR) and the California Consumer Privacy Act (CCPA), balancing personalization with user privacy is not just a technical aspiration—it is a legal and ethical necessity (Zuboff, 2019; Voigt & Von dem Bussche, 2019).

**Privacy - Preserving AI** In hyper - personalized online media centers on developing methods that minimize or eliminate the need for raw user data to be centrally collected or stored. This paper focuses on two prominent approaches: **federated learning (FL)** (McMahan et al., 2017) and **encryption - based NLP systems** (Nikolaenko et al., 2013). FL allows the model to be trained locally on user devices or edge servers without uploading sensitive data to a central server. Encryption - based techniques, including secure multi - party computation (MPC) and homomorphic encryption (HE), enable computations on encrypted data, significantly reducing the risk of data leaks (Brakerski, 2012; Acar et al., 2018).

By leveraging these technologies, media platforms can preserve user privacy while maintaining high predictive accuracy and personalization quality. However, the convergence of these techniques with advanced deep learning models (e. g., GPT - 4, PaLM, T5) poses additional challenges in terms of resource constraints, communication overhead,

complex security protocols, and interpretability (Chowdhery et al., 2022; Raffel et al., 2020; OpenAI, 2023). The research community has begun to address these challenges with specialized optimizations, novel encryption schemes, and distributed model architectures (Bonawitz et al., 2019; Li et al., 2020).

This paper provides an overview of the state - of - the - art in privacy - preserving AI for hyper - personalized online media, highlighting cutting - edge solutions in federated learning and encryption - based NLP, discussing their strengths and limitations, and drawing on recent research to propose future avenues for investigation.

## 2. Background and Motivation

### 2.1 Hyper - personalization in Online Media

Hyper - personalization refers to the delivery of highly customized content, recommendations, and interactive experiences at an individual level (Tam & Ho, 2020). Unlike one - size - fits - all recommendations, hyper - personalization uses fine - grained user data—from page dwell times to click patterns and textual interactions—to adapt content on the fly (Kapoor et al., 2019). This practice has driven user retention and monetization strategies across major platforms, but the unchecked use of personal data can lead to privacy violations and backlash from users wary of surveillance (Acquisti et al., 2016).

### 2.2 Privacy Concerns and Regulatory Landscape

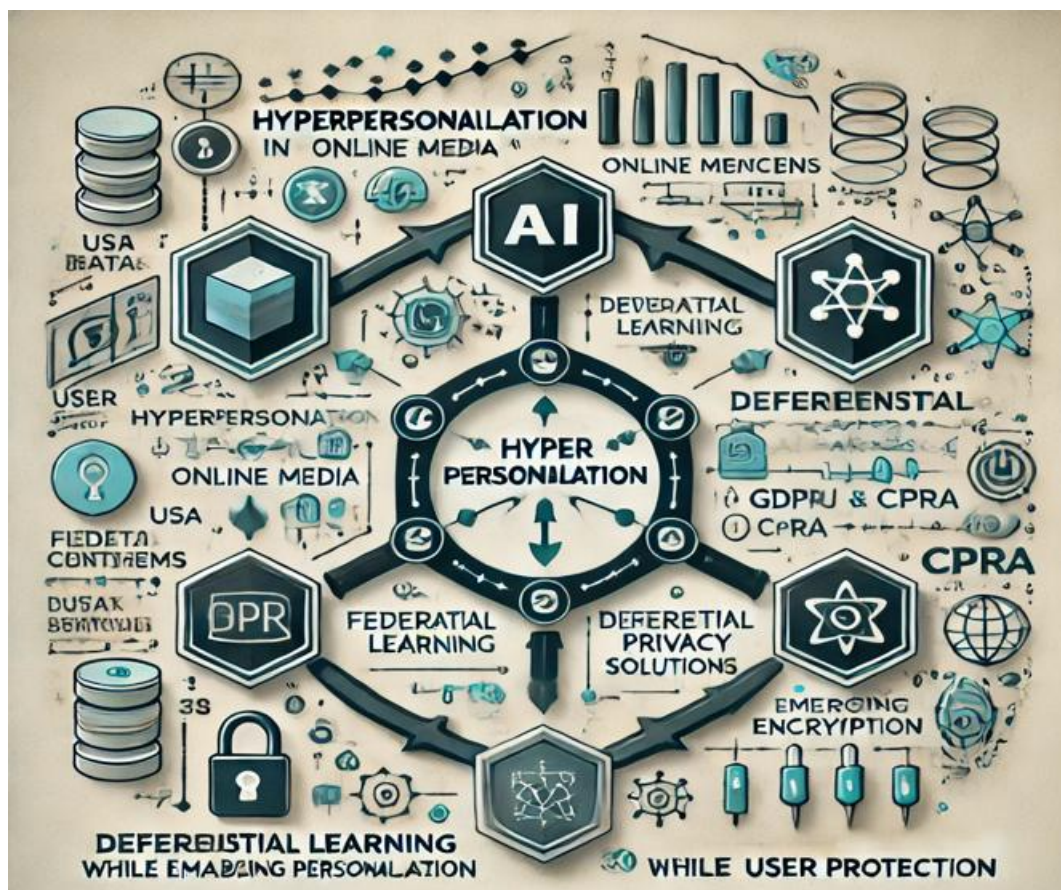
Regulatory measures such as the GDPR, the California Privacy Rights Act (CPRA), and Brazil's Lei Geral de Proteção de Dados (LGPD) impose strict requirements around data minimization, informed consent, and the right to be forgotten (Voigt & Von dem Bussche, 2019; de Lima et al., 2021). Violations carry legal and financial penalties, motivating businesses to adopt privacy - preserving techniques (Zuboff, 2019). Moreover, high - profile data breaches underscore the need for robust technical safeguards.

In hyper - personalized media contexts, any compromise of personal data—such as user preferences, conversation logs, or location data—raises significant concerns about data exploitation and user profiling (Rieke et al., 2018).

### 2.3 Emerging Privacy - Preserving AI Solutions

Privacy preservation can be broadly divided into **differential privacy** (DP) approaches (Dwork et al., 2014), **federated learning** (McMahan et al., 2017), and **encryption - based** or

cryptographic techniques (Acar et al., 2018). Federated learning gained traction for mobile keyboard suggestions in Google’s Gboard (Hard et al., 2018), while homomorphic encryption enables computations on ciphertext without revealing underlying plaintext data (Gentry, 2009). Combining these methods with advanced NLP and recommendation systems is a new frontier that seeks to maintain personalization accuracy while safeguarding user data (Liu et al., 2020; Li et al., 2020).



The diagram represents three interconnected concepts:

- 1) **Hyper - personalization:** In this section, it is explained that how some platforms use, page dwell time, the users’ click behavior, the user’s previous history to deliver a personalized message. The arrows and the icons in the dashboard presented as dynamic also show how people use real - time data to modify content in real - time.
- 2) **Privacy Concerns:** Of most importance to the discussion is the topic on the threat of personal data loss and its regulatory aspects., the shield for the protective measures in place for the users and they include symbols like GDPR and CPRA to reduce the risks involved and protect rights for the users.
- 3) **Emerging AI Solutions:** Here new trend in privacy - preserving techniques such as federated learning, differential privacy, homomorphic encryption is discussed. These approaches help in ensuring user data goes through secure processing, data

remains de - centralized, or noisy adding techniques in order to protect the data from revealing confidential information while undergoing calculations.

The movement from one section to another is well connected by thin arrows that indicate how the platform strikes a balance in targeting individual user profiles while maintaining their privacy how the ecosystem is depicted as so innovative but still adherent to ethics and law.

### 3. Federated Learning for Privacy - Preserving Personalization

#### 3.1 Federated Learning Overview

Federated learning (FL) is a collaborative machine learning paradigm that trains a shared model across multiple devices or servers holding local data samples without transferring that data to a central location (McMahan et al., 2017). The fundamental process involves each client—often a user’s device—locally training the model on its private data, then transmitting only updates or gradients to a central aggregator. The aggregator fuses these updates into a global model and sends the updated model back to all clients. This iterative process continues until the model converges to a stable state (Bonawitz et al., 2019).

### 3.2 FL for Natural Language Personalization

NLP tasks on user-generated text, such as personalized chatbots, content ranking, or writing assistants, often rely on sensitive user data (Zhang et al., 2021). FL mitigates data privacy risks by ensuring personal text never leaves the user's device. Researchers have explored FL methods for next-word prediction (Hard et al., 2018), sentiment analysis (Brisimi et al., 2018), and topic modeling (Ramage et al., 2010). This approach can be directly applied to hyper-personalized content recommendation and feed algorithms, allowing media platforms to fine-tune language models without directly accessing user logs.

### 3.3 Challenges in Federated Learning

- 1) **Communication Overhead:** FL requires iterative exchanges of model weights or gradients. Large language models exacerbate network load, especially in resource-constrained devices (Konecny et al., 2016).
- 2) **System Heterogeneity:** Users have diverse device capabilities, network connectivity, and usage patterns, leading to "straggler" issues and partial participation in training (Li et al., 2020).
- 3) **Privacy Attacks:** Although raw data is not shared, gradient-based attacks or model inversion threats can reveal sensitive information (Zhu et al., 2019). Coupling FL with differential privacy or secure aggregation can mitigate these risks (Bonawitz et al., 2017).
- 4) **Personalization vs. Global Model Trade - Off:** A global FL model may not optimally capture individual user nuances. Researchers propose personalized federated learning (PFL) to handle local preferences (Fallah et al., 2020).

### 3.4 Enhancements to Federated Learning for Hyper-personalization

- 1) **Hierarchical FL:** Organizing clients into clusters based on region or topic. Each cluster learns a specialized model, then shares cluster-level updates with a central aggregator (Brendan et al., 2021).
- 2) **On - Device Fine - Tuning:** Large pretrained models (e.g., GPT-4) can be partially fine-tuned on devices using lightweight adapter modules or low-rank adaptations, reducing the need for full model downloads (Hu et al., 2022).
- 3) **Secure Aggregation Protocols:** Ensuring encryption during gradient exchanges so the aggregator receives only masked updates (Bonawitz et al., 2017).

## 4. Encryption - Based NLP Systems

### 4.1 Homomorphic Encryption

Homomorphic encryption (HE) enables computations directly on encrypted data, preserving privacy throughout the process (Gentry, 2009). In the context of NLP personalization, a service provider can perform certain operations—like language model inference or recommendation scoring—on ciphertext user data. Users maintain control over the decryption key, ensuring that

unencrypted data never leaves their domain (Acar et al., 2018).

### 4.2 Fully vs. Partially Homomorphic Schemes

- 1) **Fully Homomorphic Encryption (FHE)** supports arbitrary computations but is often slow and resource-intensive (Brakerski, 2012).
- 2) **Partially Homomorphic Encryption (PHE)** or leveled FHE can handle certain operations (additions, multiplications) up to a specific depth, which may suffice for simpler NLP tasks (Cheon et al., 2017).

### 4.3 Secure Multi - Party Computation

Secure multi-party computation (MPC) enables a set of parties to jointly compute a function over their inputs while keeping those inputs private from each other (Yao, 1986; Lindell & Pinkas, 2009). In a personalization context, multiple stakeholders—e.g., user devices, third-party data providers, and platform servers—can collaborate on training or inference tasks without exposing raw data. Protocols like Secret Sharing or Garbled Circuits ensure data confidentiality throughout the process (Kamara et al., 2012).

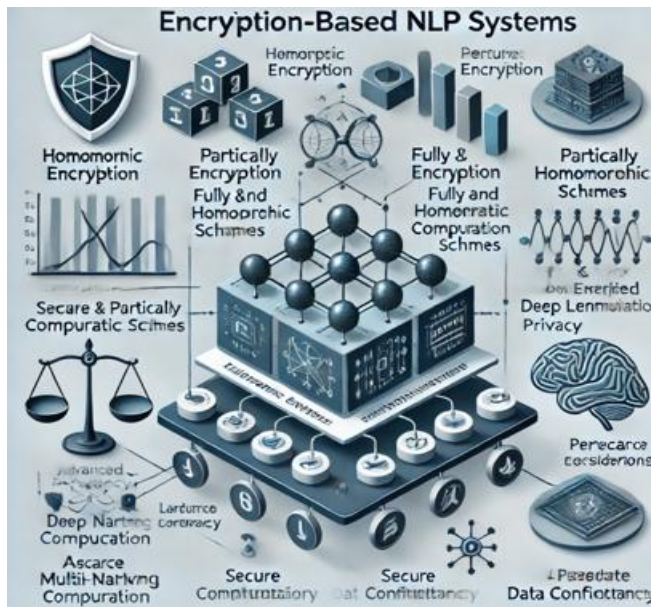
### 4.4 Combining Encryption with Advanced NLP

Modern NLP tasks rely on deep architectures with millions or billions of parameters (Devlin et al., 2019; Radford et al., 2019). Implementing these architectures on encrypted data can be prohibitively complex (Bourse et al., 2018). Recent research focuses on **model pruning, quantization, or low-rank approximations** to reduce computational overhead in encrypted domains (Falzon et al., 2022; Shao et al., 2023). Additionally, specialized frameworks like **MPC - based Transformers** attempt to replicate attention mechanisms securely (Mohassel & Rindal, 2018).

### 4.5 Performance and Practical Considerations

- 1) **Latency and Throughput:** Encryption-based operations can be 10–1000x slower than cleartext operations (Acar et al., 2018).
- 2) **Key Management:** Users must manage encryption keys, and a compromise can expose data.
- 3) **Deployment Complexity:** The orchestration of secure computation among potentially millions of user devices requires robust networking, flexible cryptographic libraries, and fallback strategies.

Despite these challenges, encryption-based NLP systems remain one of the strongest guarantees of data confidentiality, providing a path to personalization without direct data exposure (Halevi & Shoup, 2020).



Here's a brief discussion about the points represented in the diagram:

- 1) Homomorphic Encryption:** The shield in the diagram means good protection of the encrypted data. Homomorphic encryption allows computations on this data without the need to decrypt it thereby preserving user privacy when evaluating these results.
- 2) Fully vs. Partially Homomorphic Schemes:** The two superimposed icons with layers to represent addition and multiplication to show the difference between FHE, handling the far more advanced operation and PHE, which takes time but can only perform a smattering of operations.
- 3) Secure Multi - Party Computation (MPC):** The three devices interlinked by encrypted links underline the possibility of the cooperation of the stakeholders (users, platforms, and data providers) with preserving the confidentiality of data. Such secure computations are possible with help of similar protocols as Garbled Circuits and Secret Sharing.
- 4) Advanced NLP with Encryption:** Another practical example of deep learning on encrypted data is presented inside a lock by a neural network model. On performance aspect, it is found that federated learning methods such as model pruning and quantization can be considered despite it being computationally expensive.
- 5) Performance and Practical Considerations:** A balance scale represents a trade - off between latency and confidentiality. Encryption definitely introduces overhead, but nobody questions data security, this is especially important for personalized systems, and for NLP in particular.

Thus this diagram highlights the twin of data privacy and advanced NLP technologies and how hard and possible it is to create great secure systems.

## 5. Balancing Personalization and Privacy: A Comparative Analysis

### 5.1 Accuracy vs. Privacy

While FL keeps data on the client side, it can still leak patterns via gradients, necessitating additional techniques like

differential privacy, which may degrade model accuracy (Abadi et al., 2016). Encryption - based methods may impose computational overhead and limit the complexity of the model. Thus, there is a trade - off between achieving state - of - the - art hyper - personalization and maintaining strong privacy guarantees (Herlant et al., 2022).

### 5.2 Scalability and Resource Constraints

Federated learning requires iterative communication, while encryption - based approaches demand substantial computational power for cryptographic operations (Nikolaenko et al., 2013). Both approaches may strain edge devices like smartphones or IoT sensors (Burse et al., 2022). Techniques such as on - device hardware accelerators (e. g., Apple's Neural Engine) and 5G networking can partially mitigate these concerns (Park et al., 2019).

### 5.3 Interpretability and Compliance

Regulations like the GDPR's "Right to Explanation" may require businesses to explain AI - driven decisions (Goodman & Flaxman, 2017). Privacy - preserving AI methods may limit model debugging and interpretability, as the data remains encrypted or distributed (Tomsett et al., 2018). Future solutions must ensure compliance while adhering to privacy principles.

### 5.4 Ethical Implications

Balancing personalization with privacy protects user autonomy and reduces risks of digital surveillance (Zuboff, 2019). It also fosters trust in the platform, which is critical for user engagement. However, ethical dilemmas arise when personalized content can manipulate user behavior or amplify filter bubbles (Baeza - Yates, 2018). Ensuring user consent and transparency in privacy - preserving systems is fundamental to sustainable AI (Whittlestone et al., 2019).

## 6. Future Directions

### 6.1 Hybrid Approaches

When federated learning is coupled with cryptographic protocols, also known as federated analytics with secure aggregation, there is a way to address improved privacy while maintaining high model accuracy. For example, layered encryption methods can guarantee that gradients during FL updates are safe; and this minimizes inference attacks. FL has been explored to include SMPC and homomorphic encryption, which has been identified to enhance the privacy level of the learning process (Zhang et al., 2021).

### 6.2 Enhanced Cryptographic Methods

Modern cryptographic innovations provide robust frameworks for data privacy in AI applications:

- 1) Functional Encryption:** This approach allows decrypting solely some functions of the encrypted data, which should help the AI models learn different things about particular users or groups of users without being able to see their actual data. This approach is especially

beneficial to avoid the leakage of privacy - sensitive personal information (Garg et al., 2016).

- 2) **Zero - Knowledge Proofs (ZKPs):** Such cryptographic techniques are used to attest computations without revealing the input data which makes it ideal for privacy - preserving user updates or recommendation systems (Ben - Sasson et al., 2018). For instance, how do PHE guarantee that user preferences impact model changes without revealing identifiers

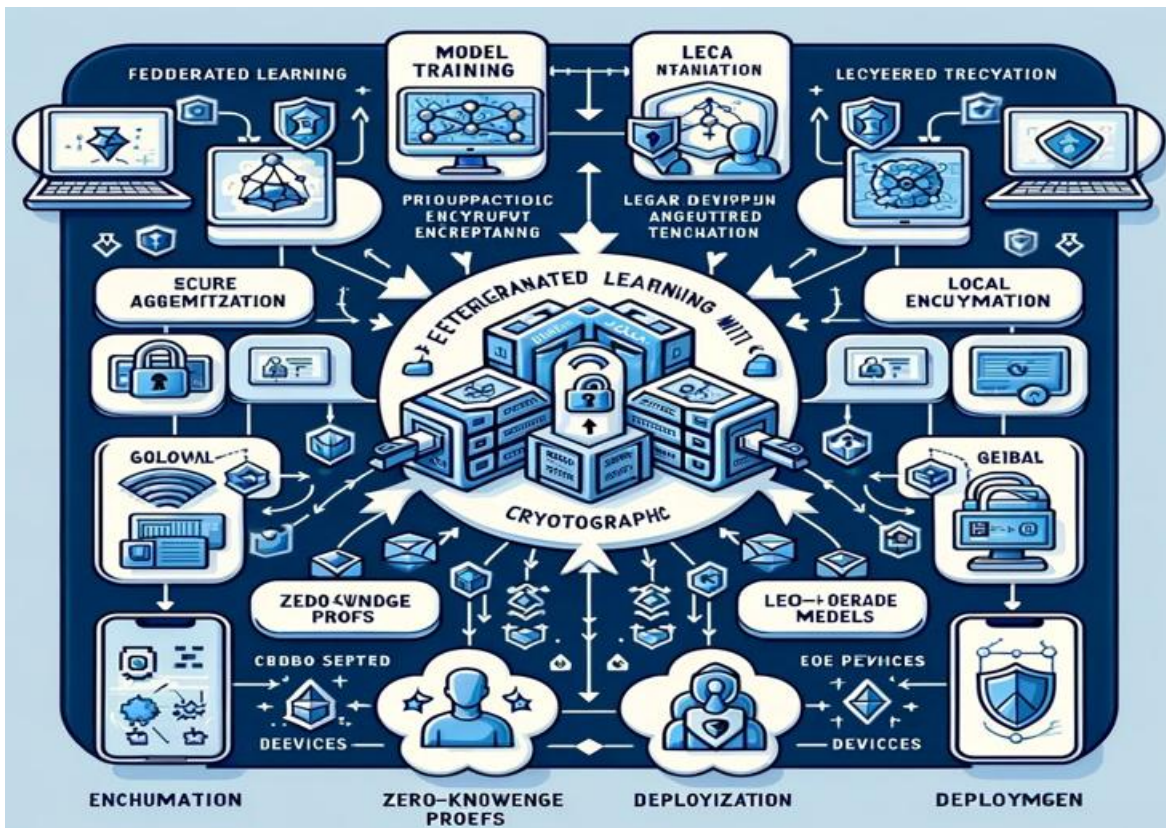
Technique	Strengths	Challenges	Applications
Federated Analytics	Combines cryptography with FL for privacy	High computational overhead	Healthcare data analysis
Layered Encryption	Protects gradients during updates	Requires advanced key management	Real - time FL model training
Zero - Knowledge Proofs	Verifies updates without exposing data	High computational costs	Secure user personalization
Model Compression	Reduces complexity for mobile applications	Trade - offs between accuracy and simplicity	On - device AI performance optimization
Personalized Federated Fine - Tuning	Balances global and local objectives	Complex privacy - accuracy trade - offs	User - specific AI customization

Model quantization has been acknowledged as one approach to model compression and acceleration.

The desire to scale up privacy - preserving AI in mobile and edge devices means that the model needs to be less complex without affecting its performance. Other methods, such as quantization, pruning, and knowledge distillation, help optimize large models for deployment, particularly in resource - limited areas (Ganesh et al., 2021). Moreover, homomorphic hardware accelerators such as the ones for HE and SMPC reduce computational burdens, allowing for real - time encrypted AI models (Dai et al., 2022).

First, we introduce a novel form of learning called personalized federated fine - tuning. Proper FedAvg and recent iterations like FedPer and FedNova try to achieve the goals of model customization locally and global model updates. These methods focus on fine - tuning the g Granted, the global shared model is kept mostly

intact for overall robustness (Hanzely and Richtárik, 2020; Deng et al., 2020). More specifically there are other more sophisticated versions of these optimizers, including differential privacy or secure gradient computations, that build upon these methods while improving their scalability and privacy (Kulkarni et al., 2023).



The diagram highlights the process of integrating federated learning (FL) with cryptographic techniques for enhanced privacy and performance:

- 1) **Data Sources:** Personal devices of users like the smartphone or computer feature an encrypted local cache, which means that no data will leave the user's environment.

- 2) **Federated Model Training:** Detailed distributed training setup have a central server for model updates while layered encryption for gradient protection during training, Thus reducing possible inference attacks
- 3) **Cryptographic Enhancements:** Functional encryption and zero - knowledge proofs are two key modules here to support computation with no disclosure of the data involved. These technologies help avoid the extraction of

unnecessary data, and raw data from the analyzed facts are not revealed.

- 4) **Global and Local Optimization:** This FL server combines global model updates with personal fine-tuning thus achieving a good trade-off between individualism and general performance.
- 5) **Deployment:** Successively, optimized and compressed models are sent to edge devices to provide highly effective and secure real-time AI applications.

The diagram visually makes the connections between such technicalities to support the main message of how the chasm between cryptography and federated learning has been bridged in the current intelligent systems.

## 7. Conclusion

Privacy in hyper-personalized AI for online media is a rapidly expanding field, driven by the need for personalized advertising and user-oriented strategies. However, this growth comes with significant challenges, including data leakage, unauthorized profiling, and user identification. Solutions like federated learning and encryption-based systems have emerged lately to address these issues by enabling AI to work without accessing raw user data.

Federated learning minimizes vulnerabilities by keeping data on users' devices while enabling AI models to improve through distributed learning. For example, language models can learn preferences without accessing actual user data, balancing privacy and personalization. Similarly, encryption techniques like homomorphic encryption and secure multiparty computation enable AI systems to operate on encrypted data, ensuring privacy while maintaining personalized services.

Despite these advancements, challenges persist. Maintaining consistent AI performance across diverse devices, such as smartphones and IoT systems, is technically demanding. Additionally, threats like adversarial attacks and privacy leaks through model gradients require robust safeguards, such as multimodal cryptographic techniques and rigorous testing. Regulatory frameworks like GDPR and CCPA further make implementation tricky, requiring a blend of legal and technical expertise to navigate compliance.

The future of privacy-preserving hyper-personalization depends on integrating advanced cryptography, distributed AI structures, and strict regulatory compliance. Future solutions must enhance efficiency, adapt to changing user and regulatory demands, and incorporate innovations like federated analytics frameworks, revolutionary cryptographic methods, and AI hardware facilitators for real-time computations on encrypted data.

Collaboration between researchers, industry leaders, and policymakers is critical. Developers must innovate in cryptography and distributed AI, while policymakers must establish frameworks that balance experimentation along with user protection. User-centered design principles must ensure that privacy-preserving technologies are intuitive, secure, and seamlessly integrated into the platforms.

If achieved, privacy-preserving AI will redefine the digital world, offering users immersive, personalized experiences

without compromising data security. It will set a new ethical standard for AI in the digital era, protecting user rights while unlocking the entire potential of personalization.

## References

- [1] Abadi, M., Chu, A., Goodfellow, I., et al. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
- [2] Acar, A., Aksu, H., et al. (2018). A Survey on Homomorphic Encryption Schemes: Theory and Implementation. *ACM Computing Surveys*, 51 (4), 1–35.
- [3] Acquisti, A., Brandimarte, L., & Loewenstein, G. (2016). Privacy and human behavior in the age of information. *Science*, 347 (6221), 509–514.
- [4] Baeza-Yates, R. (2018). Bias on the web. *Communications of the ACM*, 61 (6), 54–61.
- [5] Ben-Sasson, E., Chiesa, A., Gabizon, A., et al. (2018). Scalable Zero-Knowledge via Cycles of Elliptic Curves. *Algorithmica*, 80, 112–142.
- [6] Bonawitz, K., Ivanov, V., Kreuter, B., et al. (2017). Practical Secure Aggregation for Privacy-Preserving Machine Learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.
- [7] Bonawitz, K., Eichner, H., et al. (2019). Towards federated learning at scale: System design. *Proceedings of the 2nd SysML Conference*, 1–15.
- [8] Bourse, F., Minelli, M., Minihold, E., & Paillier, P. (2018). Fast homomorphic evaluation of deep discretized neural networks. *CRYPTO*, 483–512.
- [9] Brakerski, Z. (2012). Fully homomorphic encryption without modulus switching from classical gapsvp. *CRYPTO*, 868–886.
- [10] Brendan, M., Xie, P., et al. (2021). HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients. *ICLR*.
- [11] Brisimi, T. S., Chen, R., Mela, T., et al. (2018). Federated learning of predictive models from federated electronic health records. *International Journal of Medical Informatics*, 112, 59–67.
- [12] Burse, K., Sinha, R., & Mukhopadhyay, D. (2022). Efficient federated learning on low-resource IoT devices. *ACM Transactions on Internet of Things*, 3 (4), 1–25.
- [13] Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. *ASIACRYPT*, 409–437.
- [14] Chowdhery, A., Narang, S., Devlin, J., et al. (2022). PaLM: Scaling Language Modeling with Pathways. *arXiv preprint arXiv: 2204.02311*.
- [15] Dai, W., He, X., & Rane, S. (2022). Hardware Accelerators for Homomorphic Encryption: A Survey. *IEEE Transactions on Computers*, 71 (10), 2299–2321.
- [16] de Lima, C. H., Versieux, D., & Coelho, L. (2021). LGPD: Brazil's new data protection law. *Computer Law & Security Review*, 41, 105547.
- [17] Deng, Y., Wang, X., et al. (2020). Adaptive personalized federated learning. *arXiv preprint arXiv: 2003.13461*.

- [18] Devlin, J., Chang, M. - W., Lee, K., & Toutanova, K. (2019). BERT: Pre - training of Deep Bidirectional Transformers for Language Understanding. *NAACL - HLT*, 4171–4186.
- [19] Dwork, C., Roth, A., et al. (2014). The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9 (3–4), 211–407.
- [20] Fallah, A., Mokhtari, A., & Ozdaglar, A. (2020). Personalized Federated Learning with Theoretical Guarantees: A Model - Agnostic Meta - Learning Approach. *NeurIPS*, 3557–3568.
- [21] Falzon, S., Lal, S., & Teoh, E. (2022). Privacy - preserving inference in neural networks via model pruning and homomorphic encryption. *Information Sciences*, 601, 267–283.
- [22] Ganesh, S., Anastasopoulos, A., et al. (2021). Compressing large - scale transformer - based models: A case study on BERT. *ACL Workshops*, 15–25.
- [23] Garg, S., Gentry, C., Halevi, S., et al. (2016). Functional Encryption: Definitions and Challenges. *Information Security and Cryptology*, 325–345.
- [24] Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. *STOC*, 169–178.
- [25] Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision - making and a “right to explanation.” *AI Magazine*, 38 (3), 50–57.
- [26] Goryczka, S., & Xiong, L. (2017). A comprehensive comparison of multiparty secure additions with differential privacy. *IEEE Transactions on Dependable and Secure Computing*, 14 (5), 463–477.
- [27] Hard, A., Rao, K., Mathews, R., et al. (2018). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv: 1811.03604*.
- [28] Halevi, S., & Shoup, V. (2020). Faster homomorphic linear transformations in HELib. *CRYPTO*, 93–120.
- [29] Hanzely, F., & Richtárik, P. (2020). Federated learning of a mixture of global and local models. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 121–130.
- [30] Herlant, L., Winkler, G., & Raybourn, E. M. (2022). Impact of privacy - preserving approaches on AI performance: A systematic review. *Computers & Security*, 119, 102746.
- [31] Hu, E. J., Shen, Y., Wallis, P., et al. (2022). LoRA: Low - Rank Adaptation of Large Language Models. *ICLR*.
- [32] IEEE (2022). IEEE P2863 - Recommended Practice for Model Portability in Federated Learning. *IEEE Standards Association*.
- [33] Kamara, S., Mohassel, P., & Raykova, M. (2012). Outsourcing multi - party computation. *IACR Cryptology ePrint Archive*, 2012 (22).
- [34] Kapoor, K. K., Tamilmani, K., Rana, N. P., et al. (2019). Advances in social media research: Past, present and future. *Information Systems Frontiers*, 21 (5), 995–1011.
- [35] Konecny, J., McMahan, B., et al. (2016). Federated Optimization: Distributed Machine Learning for On - Device Intelligence. *arXiv preprint arXiv: 1610.02527*.
- [36] Krämer, J., & Böhrs, S. (2017). Let me be your personal shopper: Individualized product recommendations on the internet. *Business & Information Systems Engineering*, 59 (1), 35–45.
- [37] Kulkarni, V., Yogatama, D., & Kong, L. (2023). Towards federated hyperparameter optimization with privacy guarantees. *ICML*.
- [38] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37 (3), 50–60.
- [39] Lindell, Y., & Pinkas, B. (2009). Secure multiparty computation for privacy - preserving data mining. *Journal of Privacy and Confidentiality*, 1 (1), 59–98.
- [40] Liu, D. C., Chen, X., Meng, G., et al. (2020). Privacy - preserving deep learning for enabling big data in intelligence. *IEEE Transactions on Big Data*, 1–15.
- [41] McMahan, H. B., Moore, E., Ramage, D., et al. (2017). Communication - efficient learning of deep networks from decentralized data. *AISTATS*, 1273–1282.
- [42] Mohassel, P., & Rindal, P. (2018). ABY3: A mixed protocol framework for machine learning. *ACM CCS*, 35–52.
- [43] Nikolaenko, V., Ioannidis, S., Weinsberg, U., et al. (2013). Privacy - preserving matrix factorization. *ACM CCS*, 801–812.
- [44] OpenAI. (2023). GPT - 4 Technical Report. *arXiv preprint arXiv: 2303.08774*.
- [45] Park, J., et al. (2019). Wireless network intelligence at the edge. *Proceedings of the IEEE*, 107 (11), 2204–2239.
- [46] Radford, A., Wu, J., Amodei, D., et al. (2019). Language models are unsupervised multitask learners. *OpenAI Blog*.
- [47] Raffel, C., Shazeer, N., Roberts, A., et al. (2020). Exploring the Limits of Transfer Learning with a Unified Text - to - Text Transformer. *Journal of Machine Learning Research*, 21, 1–67.
- [48] Ramage, D., Dumais, S. T., & Liebling, D. J. (2010). Characterizing microblogs with topic models. *ICWSM*, 130–137.
- [49] Rieke, N., Hancox, J., Li, W., et al. (2018). The future of digital health with federated learning. *NPJ Digital Medicine*, 3 (119).
- [50] Shao, X., Li, Z., & Wu, Z. (2023). A framework for privacy - preserving text classification using approximate homomorphic encryption and model compression. *Information Sciences*, 623, 322–338.
- [51] Tam, C., & Ho, S. Y. (2020). Understanding the impact of personalized advertisement on consumer privacy concerns: A dual - factor model. *Information & Management*, 57 (2), 103176.
- [52] Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361 (6404), 751–752.
- [53] Tomsett, R., Braines, D., et al. (2018). Interpretable to whom? A role - based model for analyzing interpretable machine learning systems. *arXiv preprint arXiv: 1806.07552*.
- [54] Voigt, P., & Von dem Bussche, A. (2019). *The EU General Data Protection Regulation (GDPR): A Practical Guide*. Springer.
- [55] Whittlestone, J., Nyrupe, R., Alexandrova, A., & Cave, S. (2019). The role and limits of principles in AI ethics: Towards a focus on tensions. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 195–200.

- [56] Yao, A. (1986). How to generate and exchange secrets. *27th Annual Symposium on Foundations of Computer Science (FOCS)*, 162–167.
- [57] Zhang, C., Xiong, J., & Li, X. (2021). A Survey on Federated Learning for Multimedia: Communication - Efficient and Privacy - Preserving. *ACM Computing Surveys*, 54 (5), 1–36.
- [58] Zhu, L., Liu, Z., & Han, S. (2019). Deep leakage from gradients. *NeurIPS*, 14747–14756.
- [59] Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.