Federated Deep Hashing with Margin-Scalable Semantic Constraints for Privacy-Preserving Distributed Image Retrieval

Samuel D. Wells
Department of Computer Science, University of Houston, Houston, TX, USA.
wells1973@uh.edu

Dactaor Wood
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
dector.wood@uc.edu

Sean L. Bell
Department of Computer Science, George Mason University, Fairfax, VA, USA.
bellsean@gmu.edu

Summer Prasad
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
sameer.prasad792@missouri.edu

Abstract

The proliferation of image data across geographically distributed and institutionally heterogeneous repositories has created an urgent demand for retrieval systems that simultaneously guarantee high-fidelity semantic search and rigorous data privacy. Federated deep hashing, which marries the representational power of deep learning with compact binary hash codes, offers a promising pathway, yet its real-world deployment is constrained by difficulties in controlling hash code discrimination, preserving semantic alignment under non-identically distributed data, and ensuring compliance with evolving data governance regimes. This paper provides a comprehensive system-level analysis of a federated deep hashing architecture augmented with margin-scalable semantic constraints. By dynamically adjusting the similarity margins that govern hash space partitioning, the framework enables nuanced trade-offs between compactness, retrieval precision, and resilience to statistical heterogeneity across siloed data sources. We examine the structural design of federated hash learning pipelines, dissect the role of margin-scalable constraints as a governance instrument for semantic fidelity, and explore their interplay with communication efficiency, model compression, and privacy-preserving aggregation. Beyond technical architecture, the paper discusses operational infrastructure, sustainability concerns, fairness under disparate data distributions, and the alignment of federated hashing with international data protection frameworks. Through a cross-domain lens, we argue that margin scalability transforms a purely accuracy-oriented mechanism into a socio-technical control surface, empowering system operators to balance competing objectives in large-scale, privacy-sensitive image retrieval ecosystems. The analysis is grounded in contemporary advances in deep hashing, federated learning, secure aggregation, and differential privacy, providing a forward-looking perspective on resilient, policy-compliant image indexing infrastructures.

Keywords

federated learning; deep hashing; margin-scalable semantic constraints; privacy-preserving image retrieval; distributed indexing; system governance.

1. Introduction

The exponential growth of visual data generated by mobile devices, industrial sensors, biomedical imaging platforms, and surveillance networks has rendered centralized image retrieval paradigms increasingly untenable. Early approaches to content-based image retrieval relied on handcrafted local descriptors and inverted indexing schemes that, while interpretable, struggled to capture the high-level semantic concepts that human observers use to judge relevance [1]. As retrieval demands shifted toward million- and billion-scale collections, approximate nearest neighbor search techniques and product quantization methods provided necessary scalability but often at the cost of fine-grained semantic discrimination [2]. The deep learning revolution subsequently reshaped this landscape by enabling end-to-end learning of compact binary codes through deep hashing, radically compressing feature representations while preserving semantic similarity in Hamming space [3, 4, 5]. These advances, however, were designed under the implicit assumption that all training data reside in a single, centrally managed repository. In contemporary settings, images are distributed across hospitals, personal devices, retail chains, and public sector agencies, each governed by distinct privacy regulations, usage agreements, and data sovereignty requirements. This structural fragmentation has elevated the importance of federated learning as a paradigm that trains machine learning models across decentralized data without raw data exchange.

Federated deep hashing synthesizes these two trajectories, distributing the learning of hash functions across clients while aggregating model updates through a coordinating server. Despite promising demonstrations, the integration poses formidable system-level challenges. The statistical heterogeneity of client data—known as non-IIDness—can cause locally learned hash mappings to drift apart, eroding the global consistency of the semantic indexing space. Communication overhead, adversarial robustness, fairness across demographic subgroups, and compliance with data protection regulations such as the General Data Protection Regulation and the Health Insurance Portability and Accountability Act further complicate deployment. Within this intricate web of requirements, the concept of margin-scalable semantic constraints emerges as a versatile architectural lever. Borrowing from metric learning and self-supervised representation learning, these constraints allow the system to adapt the penalty margins that separate semantically similar and dissimilar pairs during hash code training. When embedded in a federated loop, such adaptability can mitigate client drift, accommodate heterogeneous similarity thresholds, and serve as a governance parameter that operators tune according to domain-specific privacy-utility trade-offs. This paper undertakes a deep analytical exploration of this system, not through the lens of narrow algorithmic optimization, but as a socio-technical infrastructure whose design decisions reverberate across communication networks, energy footprints, organizational policies, and user trust. We structure the discussion around the architectural, operational, and policy dimensions, providing a continuous academic narrative that bridges distributed systems, information retrieval, and responsible artificial intelligence.

2. Foundational Concepts and Distributed Retrieval Challenges

To understand the role of margin-scalable constraints in federated hashing, it is essential to first situate the problem within the longer arc of image retrieval and distributed computing.

Early retrieval systems treated images as bags of quantized local features, enabling text-like inverted indexes that supported real-time search over millions of items [1]. The subsequent turn toward approximate nearest neighbor search, particularly through methods based on randomized trees, locality-sensitive hashing, and product quantization, addressed the computational bottleneck of exhaustive distance calculation. Muja and Lowe's comparative study demonstrated that these techniques could dramatically accelerate retrieval, but their performance was heavily contingent on the underlying feature representations, which remained handcrafted [2]. The advent of supervised deep hashing altered the paradigm by learning feature extraction and hash coding jointly, yielding compact binary codes whose Hamming distances correlated strongly with semantic similarity. Foundational works such as spectral hashing established the principle of graph-based code learning [3], while later deep architectures like CNNH and DNNH introduced end-to-end trainable pipelines with pairwise or triplet ranking losses that enforced relative similarity constraints [4, 5]. These methods achieved remarkable accuracy on benchmark datasets such as CIFAR-10 and NUS-WIDE, yet they implicitly assumed that training and query data originated from the same distribution and resided in a single location.

The real-world deployment of such systems in distributed environments exposes several cracks. First, the assumption of central data aggregation collides with privacy regulations that prohibit transferring sensitive images outside institutional boundaries. Second, the computational and memory footprints of deep hash models, though smaller than those of full-dimensional descriptors, still demand careful management on resource-constrained edge devices. Third, the semantic boundaries learned by a globally trained hash function may inadequately capture domain-specific nuance across clients, such as differing radiological image interpretations across hospitals or varying visual search intents in cross-border e-commerce. Federated learning addresses the privacy concern by exchanging model updates rather than raw data, but introduces its own complexities: client drift due to heterogeneous local optima, straggler mitigation, and the need for secure aggregation to prevent gradient leakage that could reconstruct private images [6, 7]. These challenges form the substrate upon which federated deep hashing must be designed, and against which the efficacy of margin-scalable constraints must be evaluated.

3. Federated Learning for Decentralized Image Representation

Federated learning, as instantiated by the federated averaging algorithm and its variants, provides a general template for collaborative model training across clients holding local datasets that cannot be centralized [8]. In the context of image retrieval, this paradigm translates into the collaborative refinement of a hash function that maps raw pixels to compact binary vectors. Each client performs several local optimization steps on its private image collection and sends updated model weights or gradients to a central server, which aggregates them into a new global model. The communication-efficient nature of the protocol is critical because clients such as smartphones, rural clinics, or remote industrial cameras often operate under bandwidth constraints and intermittent connectivity. Over-the-air federated learning and hierarchical aggregation topologies further compress the communication overhead by exploiting edge servers or analog signal superposition [9, 10].

The adoption of federated learning for image retrieval is not simply a matter of replacing a centralized data loader with a federated protocol. The statistical heterogeneity of image distributions leads to divergent local hash spaces that, when naively averaged, collapse the semantic structure. A hash function learned predominantly from images of vehicles in one city may encode sedan-like shapes in a region of the binary space that conflicts with the encoding of industrial machinery from another client, resulting in confused retrieval outputs. Strategies to mitigate this drift include proximal regularization terms that tether local models to the global consensus, knowledge distillation between peers, and personalization layers that allow local deviations while preserving shared low-level features. These approaches intersect with the broader federated learning research agenda documented in comprehensive surveys [11]. Crucially, the semantic margin emerges as an additional control point: if the global aggregation imposes a rigid similarity threshold, clients with narrower intra-class variance may suffer degradation, while a margin that is too lax can blur category boundaries. Margin scalability, therefore, becomes a key enabler of robust federated image representation.

4. Deep Hashing: Compact Codes and Semantic Richness

Deep hashing has evolved from early attempts to binarize neural network activations toward highly engineered architectures that embed semantic hierarchies into compact codes. Beyond the foundational supervised methods, numerous variants have refined the learning objectives. Ranking-based deep hashing frameworks optimize a triplet loss that enforces a margin between the distances of similar and dissimilar pairs, producing codes that preserve fine-grained similarities [12]. Unsupervised and self-supervised approaches, such as topic hypergraph hashing, leverage data-internal structures to generate hash functions without requiring exhaustive pairwise labels, thereby reducing annotation costs and expanding applicability to uncurated distributed datasets [13]. The integration of attention mechanisms, generative adversarial components, and knowledge distillation has further pushed the precision-compactness frontier [14]. However, most of these advances assume a controlled training environment where hyperparameters such as the margin are set globally and remain static throughout training.

In a federated context, the assumption of a static margin is particularly limiting. The margin controls how strictly the loss function penalizes violations of the desired relative-distance ordering in the hash space. A large margin forces the hash function to separate classes with pronounced gaps, which can improve robustness to noisy features but may also discard subtle intra-class variations that are critical for fine-grained retrieval tasks, such as distinguishing similar plant species or identifying slight product defects. Conversely, a small margin allows more nuanced clustering but risks increased collision among semantically distinct images under data distribution shift. The capability to scale this margin adaptively—both across clients and over the course of training—introduces a degree of semantic governance that aligns naturally with the heterogeneous nature of federated data silos. The recently proposed margin-scalable constraints framework explicitly formulates this adaptability by enabling asymmetric semantic excavation, a process that refines the discriminative capacity of hash codes without requiring symmetrical class boundaries [15]. This innovation forms a central building block for our system-level investigation.

5. Margin-Scalable Semantic Constraints as an Architectural Lever

The notion of margin scalability in deep hashing reconfigures a traditionally fixed hyperparameter into a dynamic control mechanism. Margin-scalable semantic constraints permit the learning objective to adjust the threshold that governs the separation between hash codes of semantically similar and dissimilar images, often in a layer-wise or class-dependent manner. This adjustment can be guided by the local data statistics of each federated client, the current stage of global convergence, or externally defined quality-of-service targets. By allowing the margin to scale, the system mitigates the tension between over-generalization and over-specialization that plagues federated models trained on non-IID distributions [16]. From a systems perspective, the margin acts as a governance dial that the coordinating server can modulate to enforce shared semantic anchors while permitting local customization.

The architectural integration of such constraints requires careful coordination across the federated training loop. The global model must encapsulate not only the hash function parameters but also metadata about the learned semantic boundaries, such as per-class margin ranges or aggregated statistics of local pairwise similarities. During each communication round, clients can transmit these statistics alongside model weights, enabling the server to compute a federated margin schedule that respects the diversity of local label distributions. This schedule can be bounded by privacy-preserving mechanisms, ensuring that transmitted statistics do not leak sensitive class frequencies. The interplay between margin scalability and other federated optimization techniques, such as adaptive client weighting or momentum-based aggregation, demands a holistic design. For instance, if the server detects severe client drift through increased margin variances, it can temporarily increase proximal regularization strength or reduce the number of local epochs. Such dynamic coupling transforms the federated deep hashing system into an adaptive infrastructure that is responsive to shifting data landscapes, device availability, and retrieval accuracy requirements.

Equally important is the role of the margin in fairness. Without scalable constraints, minority classes that appear in only a few clients may be encoded with weaker discriminative power because their margin requirements differ from those of majority classes. Margin-scaling strategies that allocate larger margins to underrepresented semantic categories can partially compensate for this imbalance, promoting more equitable retrieval performance across demographic or geographic subgroups. This perspective reframes the margin not merely as a learning hyperparameter but as a policy instrument for embedding fairness objectives into the core retrieval infrastructure.

6. System Architecture: Federated Aggregation of Hash Functions

Translating federated deep hashing with margin-scalable constraints into a production-grade system necessitates a layered architecture that separates the hash function learning pipeline from the indexing and query-serving subsystems. At the top layer, a federated orchestration server manages client registration, model versioning, secure aggregation, and the dissemination of global hash models. Clients maintain local repositories of images and execute mini-batch training over private data, computing forward and backward passes through a shared backbone network, often a deep residual architecture like ResNet, followed by a fully connected hashing layer with a sign activation for binary code generation [17]. The margin-scalable loss function is computed locally, drawing on global margin schedules communicated by the server. Secure aggregation protocols, underpinned by techniques such as Shamir secret sharing or pairwise masking, ensure that individual client updates remain concealed from the server and from other clients [6]. The aggregation mechanism must operate over both the network weights and the aggregated margin-related statistics, which may include average intra-class distances or confusion matrices. These statistics are quantized and compressed before transmission to keep communication budgets within operational limits.

The indexing layer builds a distributed hash table that maps binary codes to image identifiers and optionally to the client shards where the original images reside. To support retrieval under client autonomy, each client may maintain a local inverted index of its own hash codes, while a federated search coordinator routes queries by broadcasting query hash codes and collecting ranked results. This architectural choice balances the benefits of global search against communication and privacy costs: clients that must keep image payloads local can serve results without exposing raw data, relying solely on Hamming distance comparisons. The margin scalability again plays a role here; the coordinator can apply query-time margin thresholds to filter results from clients whose semantic boundaries differ excessively from the query domain, thereby improving precision at a minor cost to recall. Such adaptive thresholding exemplifies how design-time decisions in the learning phase propagate into the operational semantics of the retrieval system.

On the infrastructure level, the architecture must accommodate a heterogeneous device landscape. Lightweight clients, such as IoT cameras, may employ model compression techniques to reduce the footprint of the hash function, using pruning, quantization, or knowledge distillation to a compact student network while participating in the federated round [14]. Compression must be orchestrated in concert with margin constraints, as aggressive pruning can inadvertently distort the margin-sensitive decision boundaries. A principled co-design of compression and margin scheduling thus becomes an essential component of the system's sustainability, reducing energy consumption and enabling deployment on ultra-low-power devices.

7. Infrastructure, Deployment, and Resource Governance

Deploying federated deep hashing at scale demands a critical examination of the infrastructure that underpins the compute, storage, and networking layers. Unlike centralized retrieval engines that can be housed in homogeneous cloud data centers, federated systems span a continuum from powerful hospital servers to battery-constrained mobile phones. This heterogeneity necessitates a flexible resource governance model that accounts for differential bandwidth, computational throughput, and availability schedules. The training orchestration must be robust to client dropout, a ubiquitous phenomenon in cross-device federated learning, and must incorporate asynchronous update mechanisms where late-arriving contributions can still be integrated without stalling the global model. Strategies such as tiered aggregation, where edge servers perform intermediate model merging before transmitting to a central cloud aggregator, can reduce wide-area network traffic and improve latency. Such infrastructure designs find corroboration in federated learning systems that exploit over-the-air computation and hierarchical architectures [9, 11].

Sustainability considerations, though often overlooked in algorithm-centric papers, are integral to the long-term viability of federated retrieval. Training deep models on thousands of edge devices for extended periods incurs a substantial carbon footprint, dominated by the electricity consumption of local training. Margin scalability can be leveraged as an efficiency lever: by dynamically adjusting margins to reduce the number of hard negative mining operations or to accelerate convergence, the total computational load can be curtailed without a commensurate loss in accuracy. Moreover, the adoption of once-for-all network designs trained jointly with a margin-focused distillation objective allows a single global model to serve diverse device profiles through sub-network extraction, minimizing redundant training cycles.

Governance of these resources extends into organizational coordination. Federated hashing deployments crossing institutional boundaries, such as a consortium of museums collaborating on a shared art retrieval system, require clear service-level agreements that specify client participation duties, data usage restrictions, and liability attribution for retrieval failures. The margin parameter, as a tunable semantic dial, can be codified into these agreements as a parameter whose range is negotiated collectively, thereby formalizing a previously ad-hoc technical choice into a governed interface. This operational framing underscores the transformation of federated deep hashing from a research prototype into a managed, sustainable infrastructure.

8. Robustness, Fairness, and Socio-Technical Resilience

A retrieval system that regularly fails for specific user demographics or under adversarial conditions cannot command the trust necessary for wide adoption. Federated deep hashing inherits the robustness challenges of both deep neural networks and distributed systems. Adversarial perturbations applied to images can flip individual hash bits, causing catastrophic ranking inversions. The decentralized nature of federated learning further exposes the system to model poisoning, where a malicious client submits crafted updates to corrupt the global hash space or to implant backdoors that trigger targeted mis-retrieval. Secure aggregation with Byzantine-resilient rules and robust statistical filtering of updates offers partial protection, but margin scalability introduces a novel defense dimension. A suspicious client that proposes excessively large margin values inconsistent with its reported data distribution can be flagged through cross-validation on public benchmark data or through differential privacy auditing of its gradient updates. Thus, the margin schedule becomes an anomaly detection substrate.

Fairness in retrieval output is equally salient. When training data across clients exhibit demographic skew—for example, if certain racial groups are underrepresented in facial image collections—the resulting hash codes may encode spurious correlations that degrade retrieval accuracy for those groups. Margin scalability, applied with a deliberate equity mandate, can counteract such skew by assigning relaxed or tightened margins to specific protected attribute groupings, provided that appropriate privacy-preserving statistics can be aggregated. The introduction of differential privacy mechanisms, which inject calibrated noise into the aggregated updates, further complicates the fairness picture because the noise distribution may differentially affect minority and majority class representations [20]. Resolving this tension requires a multi-stakeholder design process where fairness metrics are specified in collaboration with affected communities and operationalized through the margin configuration interface. In this way, the system's technical architecture becomes an expression of societal values, embodying a socio-technical contract that extends beyond traditional information retrieval evaluation metrics.

9. Policy Frameworks and Regulatory Alignment

The deployment of federated deep hashing across jurisdictions demands rigorous alignment with data protection and AI governance frameworks. The European Union's General Data Protection Regulation, California's Consumer Privacy Act, and sector-specific rules such as the Health Insurance Portability and Accountability Act impose constraints on data storage, processing, and cross-border transfer. Federated learning is often touted as a privacy-preserving technology because it keeps raw data local, but the exchange of model updates is not automatically immune from re-identification risks. Gradient inversion attacks have demonstrated that individual images can be reconstructed from shared gradients under certain conditions, thereby placing federated hashing at the intersection of technical promise and legal scrutiny. The integration of differentially private training with formal privacy budgets provides a quantifiable compliance pathway; the margin-scalable loss can be natively combined with differentially private stochastic gradient descent, where the sensitivity of margin-related updates is calibrated to bound the privacy loss [20]. The design choice to decentralize aggregation through secure multiparty computation protocols further reduces the trust surface, as codified in privacy engineering guidelines [15].

Beyond data protection, the governance of automated decision-making enshrined in AI regulations, such as the proposed EU AI Act, demands transparency and accountability for retrieval outcomes. The opacity of deep hash spaces poses a challenge for explainability. Margin schedules, being more interpretable than raw neural weights, could serve as a partial transparency mechanism: a public audit trail of how semantic boundaries were negotiated across clients during training can be logged and reviewed, offering a window into the collective decision process of the federated network. Institutional governance structures, such as federated data collaboratives with independent ethics boards, can oversee the setting of margin parameters to ensure they do not encode discriminatory preferences. The policy dimension thus completes the arc, elevating margin-scalable semantic constraints from a mere algorithmic enhancement to a tool for responsible system governance.

10. Conclusion

Federated deep hashing with margin-scalable semantic constraints represents a compelling convergence of privacy engineering, distributed systems, and semantic information retrieval. Throughout this paper, we have argued that the dynamic scalability of similarity margins transcends its role as a training hyperparameter, functioning instead as a central architectural and governance mechanism that mediates the tensions among retrieval accuracy, client drift, resource efficiency, fairness, and regulatory compliance. By coupling margin schedules with secure aggregation, model compression, and adaptive orchestration, the system can gracefully reconcile the divergent demands imposed by heterogeneous data silos, device constraints, and evolving legal landscapes. We have traced the implications from low-level infrastructure design, through robustness and fairness trade-offs, to the highest levels of policy alignment. Realizing this vision demands interdisciplinary collaboration spanning machine learning, distributed computing, privacy law, and human-computer interaction, as well as the development of new evaluation benchmarks that measure system performance under realistic non-IID, cross-jurisdictional conditions. As digital image repositories continue to expand and fragment, the ability to retrieve information while preserving the sanctity of data locality will define the next generation of intelligent retrieval infrastructures. Margin-scalable semantic constraints, embedded within a carefully governed federated framework, offer a promising foundation upon which such infrastructures can be built, ensuring that technical innovation remains accountable to the societies it serves.

References

1. Sivic, J., & Zisserman, A. (2003). Video Google: A text retrieval approach to object matching in videos. Proceedings of the IEEE International Conference on Computer Vision, 1470–1477.

2. Muja, M., & Lowe, D. G. (2014). Scalable nearest neighbor algorithms for high dimensional data. IEEE Transactions on Pattern Analysis and Machine Intelligence, 36(11), 2227–2240.

3. Weiss, Y., Torralba, A., & Fergus, R. (2008). Spectral hashing. Advances in Neural Information Processing Systems, 21, 1753–1760.

4. Xia, R., Pan, Y., Lai, H., Liu, C., & Yan, S. (2014). Supervised hashing for image retrieval via image representation learning. Proceedings of the AAAI Conference on Artificial Intelligence, 28(1), 2156–2162.

5. Lai, H., Pan, Y., Liu, Y., & Yan, S. (2015). Simultaneous feature learning and hash coding with deep neural networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3270–3278.

6. Zhang, H., Shen, F., Liu, W., He, X., Luan, H., & Chua, T. S. (2016). Supervised hashing with deep neural networks. IEEE Transactions on Image Processing, 25(8), 3676–3688.

7. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., & Wu, Y. (2014). Learning fine-grained image similarity with deep ranking. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1386–1393.

8. Zhu, L., Shen, J., Xie, L., & Cheng, Z. (2017). Unsupervised topic hypergraph hashing for efficient multimedia retrieval. IEEE Transactions on Cybernetics, 47(2), 449–460.

9. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.

10. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175–1191.

11. Yang, K., Jiang, T., Shi, Y., & Ding, Z. (2020). Federated learning via over-the-air computation. IEEE Transactions on Wireless Communications, 19(3), 2022–2035.

12. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

13. Liu, Y., Ma, Z., Liu, X., Ma, S., & Ren, J. (2019). Privacy-preserving object detection for images using edge-cloud collaborative intelligence. IEEE Internet of Things Journal, 7(5), 3800–3810.

14. Cheng, Y., Wang, D., Zhou, P., & Zhang, T. (2018). A survey of model compression and acceleration for deep neural networks. IEEE Signal Processing Magazine, 35(1), 126–136.

15. Mohassel, P., & Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. 2017 IEEE Symposium on Security and Privacy (SP), 19–38.

16. Yu, Z., Wu, S., Dou, Z., & Bakker, E. M. (2022). Deep hashing with self-supervised asymmetric semantic excavation and margin-scalable constraint. Neurocomputing, 483, 87-104.

17. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778.

18. Liu, H., Cao, J., Lu, H., Li, B., & Liu, J. (2021). Federated deep hashing for cross-modal retrieval. IEEE Transactions on Circuits and Systems for Video Technology, 32(6), 3398–3411.

19. Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. International Conference on Machine Learning, 1597–1607.

20. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.