Received: 10 July 2026; **Revised:** 10 October 2026; **Accepted:** 08 November 2026; **Published:** 20 November 2026
Archs Sci. (2026) Volume 76, Issue 5 Pages 63-68, Paper ID 2026506.
https://doi.org/10.68304/as/76506

A Comparative Study of Secure Image Deduplication Techniques Using Hashing and Deep Learning Approaches

Rahul Shah
The ICFAI University, Sikkim, Ranka Road, Lower Sichey, Gangtok, Sikkim 737101, India
Correspondence: sahrahul77@gmail.com

Abstract The rapid growth of image data in cloud storage environments has increased the need for efficient image deduplication methods that eliminate redundant information while supporting privacy, security, and storage efficiency. This paper presents a review and reported experimental comparison of image deduplication approaches, including cryptographic hashing, perceptual hashing, encryption-based methods, proof-of-ownership mechanisms, and deep learning techniques. Existing solutions are categorized according to their security functions, computational requirements, duplicate-detection capabilities, and suitability for cloud-based applications. The comparative evaluation uses the INRIA Holidays and Oxford Buildings datasets and considers SHA-256, perceptual hashing (pHash), and MobileNetV2. According to the reported results, SHA-256 identifies exact duplicate files with 100% accuracy but does not identify transformed near-duplicates, whereas pHash provides moderate near-duplicate detection. MobileNetV2 yields reported near-duplicate detection accuracies of 93.94% and 97.93% on INRIA Holidays and Oxford Buildings, respectively. These findings suggest that deep feature representations can improve the recognition of visually similar images under the specified testing conditions, although a similarity model alone does not establish cryptographic confidentiality or storage-system security. The study therefore distinguishes the accuracy of image matching from the security properties of a complete deduplication service, identifies limitations in the reported evaluation protocol, and discusses potential directions for privacy-preserving image management in cloud and Internet of Things (IoT) environments.

Index Terms Secure image deduplication; cloud storage; perceptual hashing; deep learning; MobileNetV2; near-duplicate detection

I. Introduction

The rapid development of digital technologies, mobile devices, social media platforms, and surveillance systems has resulted in a substantial increase in the volume of images generated, transmitted, and stored each day. As cloud computing has become a widely adopted platform for multimedia storage and sharing, efficiently managing increasingly large image collections has become a significant technical and operational challenge. Duplicate and near-duplicate images increase storage requirements, consume additional communication bandwidth, complicate data management, and raise operating costs for cloud service providers [1], [2]. Because images may be copied, recompressed, resized, or otherwise modified, deduplication systems must distinguish exact byte-level identity from similarity in visual content.

Image deduplication is an established technique for reducing redundant data by identifying repeated content, storing a single representative copy when appropriate, and maintaining logical references for authorized users. Although conventional deduplication can improve storage efficiency, its use with sensitive images introduces significant confidentiality, privacy, and unauthorized-access concerns. In particular, checking whether a file has already been stored can reveal information about its content or existence if the verification process is inadequately protected. Secure image deduplication therefore combines data reduction with cryptographic protection, ownership verification, and access-control mechanisms to preserve confidentiality and integrity while providing the intended storage benefits [3]–[5].

Several approaches have been proposed to address these challenges, including convergent encryption, attribute-based encryption, proof-of-ownership protocols, perceptual hashing, and deep learning-based image analysis [2], [6], [7]. Cryptographic and ownership-verification techniques primarily address confidentiality and authorization, whereas perceptual and deep feature methods primarily address similarity detection. The two functions are related but not interchangeable. Recent developments in artificial intelligence and cloud computing have improved the capacity of retrieval and deduplication systems to identify visually related images; nevertheless, a complete secure deployment must also address access control, leakage through similarity queries, and the

risks associated with replacing distinct images by one stored copy.

This paper presents a review and experimental comparison of image deduplication techniques in cloud storage settings. The methods are organized according to their operating principles, security requirements, computational costs, and application domains. In particular, a reported evaluation compares SHA-256, perceptual hashing (pHash), and MobileNetV2 using the INRIA Holidays and Oxford Buildings datasets to examine exact and near-duplicate image identification. The results are interpreted as evidence about image-matching performance under the stated evaluation conditions, not as proof that a complete cryptographic deduplication system has been secured. The analysis also identifies information needed to reproduce the experiments, including labeling rules, similarity thresholds, query-gallery splits, and the treatment of transformed images.

Secure image deduplication has attracted considerable attention because of the growing demand for efficient, privacy-preserving storage services. Early studies addressed proof-of-ownership mechanisms and private deduplication protocols intended to prevent unauthorized claims of ownership while avoiding repeated storage of identical files [7], [8]. Subsequent research examined encryption-based solutions, including convergent encryption, attribute-based protection, and related cryptographic architectures designed to improve data confidentiality during deduplication [3]–[5]. These methods do not necessarily recognize visually similar images; their primary contribution is to make duplicate handling compatible with a defined security model.

Recent developments have investigated perceptual hashing, blockchain-assisted verification, role-based re-encryption, and erasure coding to strengthen data protection, availability, and storage efficiency [9]–[11]. In parallel, visual duplicate identification has progressed from conventional image-feature matching [12] toward learning-based representations that can offer improved discrimination between related and unrelated images [2], [13]. Such developments do not remove the fundamental trade-offs among similarity-detection accuracy, confidentiality, computational overhead, and the risk of incorrectly merging distinct images. Accordingly, a comparison of secure image deduplication techniques should examine both the security guarantees of their storage protocols and the reliability of their visual matching procedures.

Table 1 summarizes representative secure image deduplication approaches reported in the literature.

Table 1: Representative Secure Image Deduplication Approaches

Key Technique	Main Contribution
Proof of Ownership [7]	Secure verification of data ownership
Private Deduplication Protocol [8]	Privacy-preserving cloud deduplication
Confidentiality-Preserving Deduplication [14]	Improved data confidentiality
Bloom Filter-Based Verification [6]	Efficient ownership verification
Erasure Codes + Encryption [9]	Reliable and secure file sharing

II. Secure Image Deduplication Techniques

Secure image deduplication integrates redundancy elimination with security mechanisms intended to protect confidential image content and ensure that access is restricted to legitimate users. Existing approaches differ in the information they compare, the extent to which images are encrypted, and the authentication procedures they require. These design choices affect computational overhead, storage savings, and exposure to potential information leakage [15]. Exact duplicate detection and visual near-duplicate detection should therefore be evaluated as separate tasks, even when both are incorporated into one operational system. Cryptographic hash functions compute compact digests of image files and are widely used to identify byte-for-byte identical content. Algorithms such as MD5, SHA-256, and SHA-3 can support exact duplicate detection efficiently; however, their security properties are not equivalent, and MD5 is unsuitable for collision-resistant security applications. Importantly, even a minor change to image bytes generally produces a different cryptographic digest, so such methods cannot directly identify near-duplicate images that have been resized, compressed, cropped, or color-adjusted [1], [5].

Perceptual hashing methods derive compact representations from visual features rather than from the exact binary contents of an image file. Thus, images with similar perceptual characteristics may produce hashes that differ by only a small Hamming distance, permitting near-duplicate identification under an appropriately selected decision threshold. Their effectiveness depends on the transformation type and strength; cropping, rotation, and substantial content changes can be more challenging than moderate adjustments in brightness or compression. Perceptual hashing consequently offers a comparatively inexpensive visual-similarity baseline for multimedia storage and retrieval [13], [16].

Encryption-based approaches aim to protect image confidentiality while permitting storage providers to identify repeated content under specified conditions. Convergent encryption and authorized or attribute-based access-control frameworks can support deduplication without simply exposing every image to an unauthorized requester. However, deterministic identifiers and content-dependent encryption may leak whether a predictable file exists unless the protocol incorporates appropriate safeguards. These issues are particularly important in cloud environments, where the server and some clients may not be fully trusted [3]–[5].

Proof-of-ownership (PoW) mechanisms require a claimant to demonstrate possession of file content before receiving access to a previously stored copy. These mechanisms address attacks in which a malicious party attempts to claim a file using only its hash or another short identifier. Bloom filters and confidentiality-preserving verification protocols have been investigated to improve the efficiency and security of ownership checks [6], [7], [14], [17]. A proof of ownership provides a particular access-control safeguard; it does not, on its own, measure whether two different image files are visually similar.

Recent advances in deep learning have expanded the range of image features available for duplicate and near-duplicate retrieval. Convolutional neural networks (CNNs) encode visual patterns that may remain informative after resizing, compression, changes in illumination, or other transformations. Deep feature matching can therefore outperform simple hashing on suitably labeled near-duplicate retrieval benchmarks, although performance depends on the network, similarity threshold, image transformations, and evaluation protocol [2], [13]. Such gains are benefits of visual representation learning and should not be interpreted as evidence of data encryption or proof-of-ownership security.

III. Comparative Analysis of Secure Image Deduplication Approaches

Secure image deduplication systems use different techniques according to the required balance of privacy protection, storage efficiency, operational cost, and visual identification accuracy. Cryptographic methods such as SHA-256 generate deterministic fingerprints from the image file and are effective for exact duplicate recognition when digests match. They do not directly recognize visually similar files with different binary content, because even a small transformation changes the resulting digest [1], [5]. Their security benefits depend on the choice of algorithm and on the surrounding storage and authorization protocol.

Perceptual hashing addresses this limitation by deriving image signatures from visual characteristics rather than exact file contents. As a result, perceptually related images may retain similar signatures even when their binary encodings differ. A suitable Hamming-distance threshold can identify potential near-duplicate pairs with moderate processing requirements, although threshold calibration and transformation sensitivity remain important practical considerations [13], [16].

Encryption-based techniques, including convergent encryption and attribute-based access mechanisms, are designed to protect stored content while enabling duplicate detection and appropriately restricted sharing. Their principal objective is confidentiality and authorized use, rather than improved discrimination between near-duplicate visual scenes. Consequently, their security effectiveness should be established by a threat model and protocol analysis instead of being inferred from image-matching accuracy [3]–[5].

Recent research has increasingly examined deep learning-based approaches that extract image representations suitable for similarity comparisons. Such feature representations can identify visually related images despite transformations including cropping, compression, brightness adjustment, and some types of flipping. However, the degree of robustness depends on feature extraction, training conditions, and evaluation thresholds. Deep learning also requires additional computation relative to cryptographic or simple perceptual hashes, thereby introducing a practical trade-off between detection quality and processing time [2], [13].

Table 2 summarizes the major image deduplication approaches. Its qualitative security and cost labels are descriptive classifications based on the cited literature, not numerical results of a common cryptographic security assessment.

Table 2: Comparison of Major Image Deduplication Approaches

Approach	Security	Exact Duplicate Detection	Near-Duplicate Detection	Computational Cost
SHA-256 and MD5	High	High	Low	Low
Perceptual Hashing	Medium	High	Medium	Low
Convergent Encryption	High	High	Low	Medium
Attribute-Based Encryption	High	High	Low	Medium
Deep Learning (CNNs)	Medium	High	High	High

IV. Cloud Storage Solutions for Secure Image Deduplication

Cloud storage platforms provide scalable infrastructure for managing large image repositories and for processing repeated uploads across users and devices. Object storage systems can reduce redundant data by maintaining a single stored representation and associating it with separate logical references where permissible. Secure deduplication adds controls for encryption, authentication, ownership verification, integrity, and authorization so that storage savings do not create new opportunities for unauthorized disclosure [3], [4], [8].

Cloud-based deduplication can lower storage consumption and network transfer requirements by avoiding repeated transmission or storage of identical content. When multiple authorized users refer to one stored image, logical references preserve access without requiring separate physical copies. However, these efficiency benefits depend on the duplication rate, reference-management design, and the costs of indexing and verification. Encryption, proof-of-ownership protocols, and access-control policies must also be implemented consistently to prevent unauthorized access to already stored files [6], [7], [14].

Despite these potential benefits, cloud-based image deduplication faces persistent challenges in confidentiality, scalability, computational cost, and the accurate recognition of near-duplicate content in heterogeneous cloud and IoT deployments [18]. Integrating cryptographic storage mechanisms with perceptual hashes or deep features may provide better functionality across the full workflow, but it also creates additional questions about similarity-score leakage, false matches, and the secure management of references to retained images [2], [13].

V. Experimental Evaluation

To compare the image-identification performance of the selected techniques, the manuscript reports experiments based on two established image-retrieval benchmarks: INRIA Holidays, associated with the work of Jégou et al. [19], and Oxford Buildings, associated with the object-retrieval study of Philbin et al. [20]. The INRIA Holidays evaluation contains 1,491 original images, 1,491 exact duplicates, and 4,473 near-duplicates generated by brightness adjustment, cropping, and flipping, giving 7,455 test images. The Oxford Buildings evaluation contains 5,063 originals, 5,063 exact duplicates, and 5,063 transformed near-duplicates, giving 15,189 images. The two constructed evaluation sets therefore contain 22,644

reported images in total, as summarized in Table 3. These totals describe the expanded evaluation collections rather than the original sizes of both public benchmarks. The source manuscript does not identify the precise transformation parameters, duplicate-generation scripts, or data split definitions, which are needed for exact replication.

Table 3: Composition of the Image Deduplication Evaluation Sets

Dataset	Original Images	Exact Duplicates	Near Duplicates	Total Images
INRIA Holidays	1491	1491	4473	7455
Oxford Buildings	5063	5063	5063	15189
Total	6554	6554	9536	22644

Three representative image-matching approaches are compared: SHA-256, perceptual hashing (pHash), and MobileNetV2. SHA-256 is used as a byte-level cryptographic fingerprint for exact duplicate detection, whereas pHash provides a perceptual signature for comparing visually related images. The MobileNetV2 architecture, introduced by Sandler et al. [21], supplies learned image representations for cosine-similarity comparisons. These baselines answer different questions: cryptographic hashing tests file identity, while pHash and MobileNetV2 test visual similarity. None of the three methods, when evaluated without a surrounding cryptographic protocol, is sufficient to demonstrate a complete secure cloud deduplication system.

The reported experiments were performed on the Kaggle cloud platform using TensorFlow 2.19 and an NVIDIA Tesla P100 GPU. Detection accuracy and processing time were selected as the principal quantitative measures. For MobileNetV2, the input images were resized to 224 $imes$ 224 pixels, and cosine similarity was used to compare extracted feature vectors. However, the source does not specify the pretrained weight version, classification or embedding layer, similarity thresholds, matching-pair construction, or whether the time measurements include feature extraction, hashing, and gallery searching. These details are essential for independently checking the numerical results and ensuring that the processing-time comparison is fair.

The reported experimental results indicate clear differences between byte-level matching and visual near-duplicate recognition. SHA-256 correctly identifies the exact duplicates in the evaluated samples but is ineffective for identifying transformed near-duplicates because these images no longer contain the same file bytes. Perceptual hashing identifies some such pairs by comparing image characteristics; nevertheless, the reported results indicate reduced performance when the transformation involves cropping, flipping, or changes in brightness. These are observations under the reported tests and should be interpreted in relation to the matching thresholds and image transformations, which are not fully documented.

MobileNetV2 achieved the highest reported near-duplicate detection accuracy by extracting a deep visual representation rather than comparing raw file digests. Table 4 records 93.94% near-duplicate accuracy for MobileNetV2 on the INRIA Holidays evaluation, compared with 34.25% for pHash and 0% for SHA-256. The supplementary Oxford Buildings evaluation reports 4,957 correctly detected images among 5,062 tested images, corresponding to approximately 97.93% accuracy (Table 5). The difference between 5,063 near-duplicate images listed for the Oxford collection and 5,062 images tested in the validation must be explained from the experimental records. The reported accuracies should not be presented as independently reproduced measurements until the underlying code, test labels, and decision thresholds are available.

Table 4: Reported Comparison of Deduplication Methods on INRIA Holidays

Method	Exact Duplicate Accuracy (%)	Near-Duplicate Accuracy (%)	Processing Time (s)
SHA-256	100.00	0.00	11.87
pHash	100.00	34.25	107.50
MobileNetV2	100.00	93.94	178.02

Table 5: Reported MobileNetV2 Validation on Oxford Buildings

Dataset	Valid Images Tested	Correctly Detected	Accuracy (%)
Oxford Buildings	5062	4957	97.93

Figure 1 compares the reported near-duplicate detection accuracy of the evaluated approaches on the INRIA Holidays-derived test collection. MobileNetV2 outperforms SHA-256 and pHash on this visual-similarity task, although SHA-256 serves a different technical purpose: detecting exact file identity. Figure 2 and Table 4 compare processing times and illustrate the trade-off between computational cost and reported similarity-detection performance. MobileNetV2 takes 178.02 seconds, compared with 107.50 seconds for pHash and 11.87 seconds for SHA-256 under the supplied test configuration. These timing values should be compared cautiously until it is established that each method was timed over the same image set and processing stages. The reported visual matching advantage does not independently demonstrate confidentiality, encryption effectiveness, or authorized access in a cloud service.

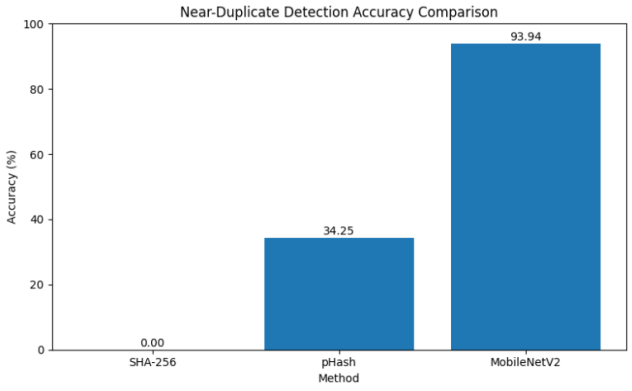


Figure 1: Near-duplicate detection accuracy comparison of SHA-256, pHash, and MobileNetV2 on the INRIA Holidays dataset.

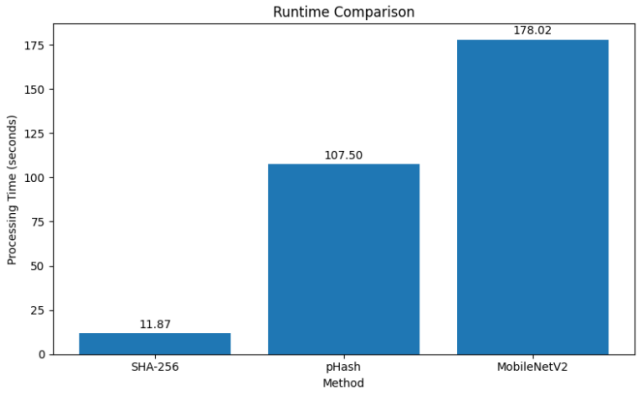


Figure 2: Processing time comparison of SHA-256, pHash, and MobileNetV2 on the INRIA Holidays dataset.

VI. Challenges

Despite considerable progress, secure image deduplication continues to face several technical challenges. Strong encryption may make content comparison difficult, whereas deterministic identifiers and similarity-search interfaces can reveal information that should remain private. Large cloud repositories also require methods capable of indexing and comparing millions of images within acceptable computational and memory budgets. Although SHA-256 accurately identifies byte-identical content, it does not directly recognize transformed near-duplicates, whereas perceptual hashes may fail under complex image edits. Similar challenges occur in surveillance and face deduplication applications [22]. Deep feature extraction improves the detection of certain visually related images but increases processing time and demands a clearly defined threshold for deciding when images should be considered duplicates. Finally, the possibility of falsely merging distinct images must be evaluated separately from the goal of maximizing storage reduction; confidentiality, integrity, and ownership assurance remain independent security requirements.

VII. Future Directions

Future research on secure image deduplication should investigate lighter-weight and privacy-preserving feature extractors that maintain near-duplicate detection performance while reducing computational costs. Integrating duplicate identification with image authenticity verification and forgery detection may further improve the integrity of multimedia repositories [23]. Federated learning and blockchain-assisted logging are potential components for decentralized or auditable services, although they introduce their own security, resource, and governance requirements. Hybrid schemes combining cryptographic identifiers, authenticated proof-of-ownership protocols, and deep visual features may offer a more practical balance between accurate matching and confidentiality. Future experiments should report thresholds, precision, recall, false-merger rates, storage savings, throughput, and tests against adversarially modified images, rather than relying on accuracy alone. These requirements are particularly important for cloud-edge and IoT systems with limited computing resources.

VIII. Conclusion

This paper presented a review and reported experimental evaluation of image deduplication techniques for cloud storage environments. The analysis considered cryptographic hashing, perceptual hashing, encryption-enabled deduplication, proof-of-ownership protocols, and deep feature-based similarity matching. It distinguished the security properties of storage and authorization methods from the image-similarity capabilities of feature extraction and matching methods. The relative merits of the approaches were discussed in terms of duplicate-detection accuracy, processing requirements, storage efficiency, and the need for confidentiality-preserving implementation.

The reported experiments used augmented sets derived from the INRIA Holidays and Oxford Buildings benchmarks. Under the stated conditions, SHA-256 achieved 100% accuracy on byte-identical copies but 0% on the evaluated near-duplicate images. Perceptual hashing improved near-duplicate matching relative to SHA-256, while MobileNetV2 achieved the highest reported visual detection rates: 93.94% on INRIA Holidays and 97.93% on Oxford Buildings. These results support a distinction between exact file matching and visual similarity detection, but they do not establish a secure storage deployment or independently verify the reported experimental outcomes. Such conclusions require a transparent evaluation protocol, reproducible test labels, and clearly defined cryptographic safeguards.

Overall, secure image deduplication can reduce unnecessary storage and data-transfer requirements when it is implemented with suitable authorization, integrity checks, and privacy safeguards. Deep feature extraction offers a promising way to recognize images that are visually similar despite differences in file encoding, while cryptographic hashing remains appropriate for byte-level duplicate detection. A complete future system should combine these complementary functions and assess not only detection accuracy but also false merges, storage savings, throughput, and resistance to information leakage. The present findings are therefore best regarded as a method comparison that motivates further end-to-end research on secure multimedia storage in cloud and IoT environments.

Conflict of Interest

The author declares that there are no known conflicts of interest that could have influenced the research synthesis, analysis, or presentation of this review.

Funding Information

The author declares that no specific financial support was received from public, commercial, or not-for-profit funding agencies for the preparation and analysis of this study.

Data Availability

The study uses the publicly available INRIA Holidays and Oxford Buildings image-retrieval benchmarks and reports evaluations using additional exact and transformed duplicate images. The specific transformed image sets, matching labels, evaluation scripts, model configuration, and execution logs were not included with the supplied manuscript. Requests for these additional materials may be directed to the author; their availability and any sharing restrictions should be confirmed before publication.

References

[1] Rashid, F., Miri, A., & Woungang, I. (2016). Secure image deduplication through image compression. *Journal of Information Security and Applications, 27–28*, 54–64. https://doi.org/10.1016/j.jisa.2015.11.003

[2] Kaur, R., Bhattacharya, J., & Chana, I. (2022). Deep CNN based online image deduplication technique for cloud storage system. *Multimedia Tools and Applications, 81*(28), 40793–40826. https://doi.org/10.1007/s11042-022-13182-7

[3] Cheng, S., Tang, Z., Zeng, S., Cui, X., & Li, T. (2024). PFDup: Practical fuzzy deduplication for encrypted multimedia data. *Journal of Industrial Information Integration, 40*, Article 100613. https://doi.org/10.1016/j.jii.2024.100613

[4] Song, M., Hua, Z., Zheng, Y., Xiang, T., & Jia, X. (2024). SimLESS: A secure deduplication system over similar data in cloud media sharing. *IEEE Transactions on Information Forensics and Security, 19*, 4700–4715. https://doi.org/10.1109/TIFS.2024.3382603

[5] Storer, M. W., Greenan, K., Long, D. D. E., & Miller, E. L. (2008). Secure data deduplication. In *Proceedings of the 4th ACM International Workshop on Storage Security and Survivability* (pp. 1–10). Association for Computing Machinery. https://doi.org/10.1145/1456469.1456471

[6] Blasco, J., Di Pietro, R., Orfila, A., & Sorniotti, A. (2014). A tunable proof of ownership scheme for deduplication using Bloom filters. In *2014 IEEE Conference on Communications and Network Security* (pp. 481–489). IEEE. https://doi.org/10.1109/CNS.2014.6997518

[7] Halevi, S., Harnik, D., Pinkas, B., & Shulman-Peleg, A. (2011). Proofs of ownership in remote storage systems. In *Proceedings of the 18th ACM Conference on Computer and Communications Security* (pp. 491–500). Association for Computing Machinery. https://doi.org/10.1145/2046707.2046765

[8] Ng, W. K., Wen, Y., & Zhu, H. (2012). Private data deduplication protocols in cloud storage. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing* (pp. 441–446). Association for Computing Machinery.

[9] Jacob, C., & Rekha, V. R. (2017). Secured and reliable file sharing system with de-duplication using erasure correction code. In *2017 International Conference on Networks & Advances in Computational Technologies (NetACT)* (pp. 221–228). IEEE.

[10] Uthayashangar, S., Abinaya, J., Harshini, V., & Jayavardhani, R. (2020). Image and text encrypted data with authorized deduplication in cloud. In *2020 International Conference on System, Computation, Automation and Networking (ICSCAN)* (pp. 1–5). IEEE.

[11] Lee, D., & Park, N. (2019). CCTV video data deduplication method using blockchain and big data analysis. *International Journal of Advanced Science and Technology, 28*(3), 267–275.

[12] Chen, M., Wang, S., & Tian, L. (2013). A high-precision duplicate image deduplication approach. *Journal of Computers, 8*(11), 2768–2775. https://doi.org/10.4304/jcp.8.11.2768-2775

[13] Zhou, Z., Wu, Q. M. J., Huang, F., & Sun, X. (2017). Fast and accurate near-duplicate image elimination for visual sensor networks. *International Journal of Distributed Sensor Networks, 13*(2), Article 1550147717694172. https://doi.org/10.1177/1550147717694172

[14] González-Manzano, L., & Orfila, A. (2015). An efficient confidentiality-preserving proof of ownership for deduplication. *Journal of Network and Computer Applications, 50*, 49–59. https://doi.org/10.1016/j.jnca.2014.12.004

[15] Usharani, S., & Dhanalakshmi, K. (2023). An image storage duplication detection method using recurrent learning for smart application services. *The Journal of Supercomputing, 79*(10), 11328–11354. https://doi.org/10.1007/s11227-023-05042-4

[16] Layek, A. K., Gupta, A., Ghosh, S., & Mandal, S. (2016). Fast near-duplicate detection from image streams on online social media during disaster events. In *2016 IEEE Annual India Conference (INDICON)* (pp. 1–6). IEEE.

[17] Di Pietro, R., & Sorniotti, A. (2012). Boosting efficiency and security in proof of ownership for deduplication. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security* (pp. 81–82). Association for Computing Machinery.

[18] Yan, H., Li, X., Wang, Y., & Jia, C. (2018). Centralized duplicate removal video storage system with privacy preservation in IoT. *Sensors, 18*(6), Article 1814. https://doi.org/10.3390/s18061814

[19] Jégou, H., Douze, M., & Schmid, C. (2008). Hamming embedding and weak geometric consistency for large-scale image search. In *Computer vision—ECCV 2008* (pp. 304–317). Springer. https://doi.org/10.1007/978-3-540-88682-2_24

[20] Philbin, J., Chum, O., Isard, M., Sivic, J., & Zisserman, A. (2007). Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1–8). IEEE. https://doi.org/10.1109/CVPR.2007.383172

[21] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510–4520). https://doi.org/10.1109/CVPR.2018.00474

[22] Chen, Q., Yang, L., Zhang, D., Shen, Y., & Huang, S. (2018). Face deduplication in video surveillance. *International Journal of Pattern Recognition and Artificial Intelligence, 32*(3), Article 1856001. https://doi.org/10.1142/S0218001418560013

[23] Kharat, J., & Chougule, S. (2020). A passive blind forgery detection technique to identify frame duplication attack. *Multimedia Tools and Applications, 79*(11–12), 8107–8123. https://doi.org/10.1007/s11042-019-08272-y