Enhancement-Driven Pretraining for Robust Fingerprint Representation Learning

Ekta Gavas[1][a], Kaustubh Olpadkar[2][b] and Anoop Namboodiri[1][c]

[1]*Centre for Visual Information Technology, International Institute of Information Technology, Hyderabad, India*

[2]*Stony Brook University, U.S.A.*

Keywords: Fingerprint Representation Learning, Fingerprint Verification, Self-Supervised Learning, Deep Learning.

Abstract: Fingerprint recognition stands as a pivotal component of biometric technology, with diverse applications from identity verification to advanced search tools. In this paper, we propose a unique method for deriving robust fingerprint representations by leveraging enhancement-based pre-training. Building on the achievements of U-Net-based fingerprint enhancement, our method employs a specialized encoder to derive representations from fingerprint images in a self-supervised manner. We further refine these representations, aiming to enhance the verification capabilities. Our experimental results, tested on publicly available fingerprint datasets, reveal a marked improvement in verification performance against established self-supervised training techniques. Our findings not only highlight the effectiveness of our method but also pave the way for potential advancements. Crucially, our research indicates that it is feasible to extract meaningful fingerprint representations from degraded images without relying on enhanced samples.

1 INTRODUCTION

Fingerprint recognition remains a cornerstone in biometric identification, valued for its uniqueness, permanence, and user-friendliness (Maltoni et al., 2022; Wayman et al., 2005; Allen et al., 2005). As demand in law enforcement, personal identification, and secure authentication continues to rise, the need to enhance precision and efficiency in fingerprint recognition systems becomes increasingly vital (Allen et al., 2005).

Despite advancements in the field, challenges persist, including handling partial or distorted fingerprints, managing high interclass similarity, and addressing the expansive dimensionality of the feature space (Maltoni et al., 2022; Hong et al., 1998; Cappelli et al., 2007). Many state-of-the-art works in fingerprint matching rely on minutia-based approaches (Ratha et al., 1996; Chang et al., 1997; Maltoni et al., 2022; Cappelli et al., 2010b; Cappelli et al., 2010a; Jain et al., 2001; Jain et al., 1997). This involves extracting minutiae and matching templates to determine similarity, but traditional minutia-based methods face limitations like noise sensitivity and difficulty with partial prints (Maltoni et al., 2009; Hong et al., 1998; Maltoni et al., 2009; Zaeri, 2011).

In contrast, Convolutional Neural Networks (CNNs) present a contemporary solution, effectively overcoming limitations and improving accuracy. CNNs handle partial prints, tolerate distortions, and adapt to diverse finger conditions, showcasing scalability and efficient comparison even with growing databases (Nguyen et al., 2018; Deshpande et al., 2020; Darlow and Rosman, 2017; Tang et al., 2017; Engelsma et al., 2019).

The surge in self-supervised learning techniques in machine learning has extended to fingerprint biometrics (Jaiswal et al., 2020; Liu et al., 2021; Jing and Tian, 2020). Offering solutions to challenges in data acquisition, self-supervised learning bypasses time-consuming labeling processes. In this paper, we explore the potential of deep CNNs for superior matching performance, proposing a pretraining technique based on U-Net for fingerprint enhancement. The U-Net model, known for biomedical image segmentation (Ronneberger et al., 2015), effectively enhances fingerprints by extracting contextual information, aiming to derive compact, discriminative fingerprint embeddings. Our study pursues two objectives: proposing a pretraining technique with

[a] https://orcid.org/0000-0001-6437-3357
[b] https://orcid.org/0009-0008-3811-4771
[c] https://orcid.org/0000-0002-4638-0833

Gavas, E., Olpadkar, K. and Namboodiri, A.
Enhancement-Driven Pretraining for Robust Fingerprint Representation Learning.
DOI: 10.5220/0012474900003660
Paper published under CC license (CC BY-NC-ND 4.0)
In *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2024) - Volume 2: VISAPP*, pages 821-828
ISBN: 978-989-758-679-8; ISSN: 2184-4321
Proceedings Copyright © 2024 by SCITEPRESS – Science and Technology Publications, Lda.

U-Net and assessing the efficacy of these representations through verification performance against existing self-supervised methods. We experiment with training and inference techniques to optimize the use of representations for fingerprint verification tasks. This paper aims to deepen our understanding of fingerprint recognition, inspiring future progress in this direction.

1.1 Contributions

Here are the main contributions of this work:

1. We suggest a pre-training technique with fingerprint enhancement task on our encoder and demonstrate the usefulness of this approach in representation learning in self-supervised setting.
2. We describe a method to fine-tune the learned embeddings for fingerprint verification task.
3. We evaluate our approach with various evaluation metrics demonstrating its effectiveness in fingerprint verification task and also provide a comparison with previous state-of-the-art self-supervised learning methods.

1.2 Related Work

The need for improved fingerprint recognition tools has spurred the development of effective fingerprint representation methods. Various approaches, drawing on domain knowledge, have enhanced the accuracy and speed of fingerprint identification (Engelsma et al., 2019; Tang et al., 2017). This paper explores a pretraining technique, focusing on an enhancement task to optimize model learning for representation.

1.2.1 Image Enhancement

Early fingerprint image enhancement methods, such as Gabor filters and Fourier Transform, faced challenges with poor quality, noise, and pattern variations (Greenberg et al., 2002; Hong et al., 1998; Kim et al., 2002; Yang et al., 2002; Liu et al., 2014; Sherlock et al., 1992; Chikkerur et al., 2005; Rahman et al., 2008). Convolutional Neural Networks (CNNs), adept at hierarchical learning, have proven effective in capturing minutiae and latent features, enhancing recognition accuracy (Nguyen et al., 2018; Deshpande et al., 2020; Tang et al., 2017). U-Net, originally designed for biomedical image segmentation, has been adapted for fingerprint enhancement (Ronneberger et al., 2015). Various modifications to U-Net, tailored for fingerprint enhancement tasks, have been proposed (Gavas and Namboodiri, 2023; Qian et al., 2019; Liu and Qian, 2020).

1.2.2 Self-supervised Learning Techniques

Self-supervised learning, an alternative to traditional supervised learning, capitalizes on unlabeled data using pretext tasks for feature representation (Jaiswal et al., 2020; Jing and Tian, 2020). Contrastive learning, a cornerstone of self-supervised learning, differentiates between similar and dissimilar instances (Liu et al., 2021). Techniques like SimCLR, MoCo, BYOL, SwAV, and Noise Contrastive Estimation showcase the diversity of contrastive learning approaches (Chen et al., 2020a; Chen et al., 2020b; Grill et al., 2020; Caron et al., 2021; Gutmann and Hyvärinen, 2010). These methods provide insights into contrastive learning's potential applications in fingerprint biometrics.

2 METHODOLOGY

The methodology for our research is constructed around a two-stage framework to probe the potential of self-supervised learning in fingerprint representation learning. A broad overview of the process is as follows:

- **Stage 1: Self-Supervised Pre-training:** This is the initial stage of our methodology, in which we perform pre-training of our models in a self-supervised manner. It includes the application of both existing self-supervised learning techniques as well as our novel enhancement-based approach for this task. This stage intends to leverage the power of unlabeled data to learn meaningful representations that can serve as a starting point for subsequent stages. Notably, for all methods, we keep the encoder architecture the same. While other self-supervised methods traditionally use encoders like ResNet or Vision Transformers, in our framework we use the encoder of our U-Net-based model to ensure a fair comparison.
- **Stage 2: Probing Experiments:** Upon completion of the pre-training phase, we progress to the second stage where a few linear layers (MLP) are added on top of the frozen pre-trained encoder, making the representations 512-d. We then perform probing experiments using this newly formed model. By keeping the encoder part frozen, we ensure that the model adapts the existing representations for the verification task without altering the learned patterns from the self-supervised pre-training phase.

Following this framework, we navigate through the process of adapting and implementing self-

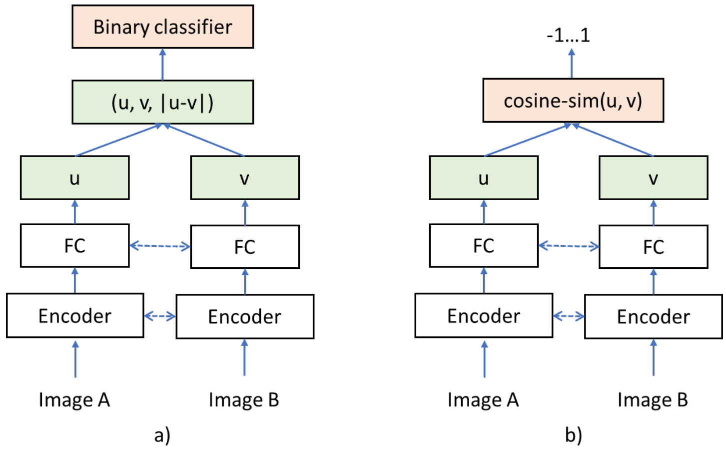

Figure 1: a) Architecture with verification objective i.e with binary classifier (at training and inference) b) Architecture to compute similarity scores (at inference). The dotted arrows indicate networks having tied weights (siamese network structure).

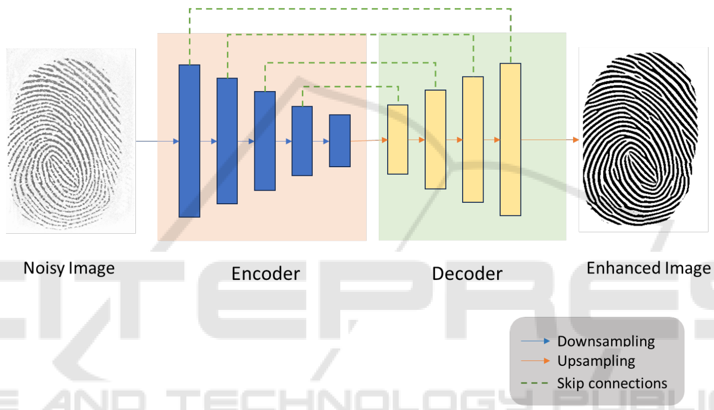

SCITEPRESS
SCIENCE AND TECHNOLOGY PUBLICATIONS

Figure 2: U-Net architecture for enhancement task for the pre-training stage in the self-supervised setting. For representation learning, the decoder is discarded and the binary classifier is attached.

supervised learning techniques, exploring a U-Net-based pre-training strategy, and conducting probing experiments with pre-trained networks. The sections below provide a detailed overview of the procedures involved in each stage.

2.1 U-Net-Based Pretraining

While applying existing self-supervised methods to fingerprint data offers a valuable starting point, we advocate for a self-supervised learning method tailored specifically for the uniqueness of fingerprint data. Drawing on our insights from U-Net-based enhancement works, our approach employs the training of a fingerprint enhancement model as a form of self-supervision.

We employ U-Net-based fingerprint enhancement for pre-training, hypothesizing that the U-Net encoder, trained on fingerprint enhancement, holds valuable fingerprint representations. Enhancing a fingerprint image becomes an effective self-supervised task, encouraging the model to learn useful, fingerprint representations. The pre-trained encoder already encapsulates crucial information about the fingerprint, providing a foundation for further representation learning. The quality of these initial representations hinges on the efficacy of the U-Net-based enhancement model, emphasizing the significance of the model's design and training.

For enhancement-based pre-training, we use the basic U-Net architecture (Figure 2) to optimize the fingerprint enhancement task. This simple image-to-image network takes a degraded fingerprint image as input, degraded with various noises. The network aims to predict an enhanced version of the fingerprint image by removing noise while maintaining and restoring the ridge structure. This ensures the network learns minute details of fingerprint structure and enhances it where possible, aiding robust feature representation extraction later. We term it self-supervision as we use supervision from the enhancement task indirectly. This leverages a smaller amount of labeled data with limited impressions and identities.

Table 1 are the results of the first stage of our network where we are pre-training the U-Net model for enhancement task. Results of this pre-training stage are demonstrated in Table 1.

2.2 Learning Fingerprint Representation

After the self-supervised pre-training, we conduct the probing experiments using the pre-trained networks. These experiments aim to assess the usefulness of the learned representations for the task of fingerprint verification. For this, we add a 3-layer MLP projection head on top of the frozen encoder part of the pre-trained network. We then train this model using a Sentence-BERT-like (Reimers and Gurevych, 2019) siamese architecture, with a limited amount of labeled data for the fingerprint verification task. We concatenate the fingerprint representations u and v of the image pair with the element-wise difference $|u-v|$ and then pass it through the linear layers and train it for binary-classification objective as illustrated in Figure 1. By keeping the encoder part frozen, the model learns to adapt the existing representations for the verification task, without changing the underlying learned patterns. This approach allows us to leverage a large amount of unlabeled data to learn initial representations and a limited amount of labeled data for supervised adaptation. Note that in the supervised fine-tuning, allowing modifications in the encoder weights can lead to higher performance on the end task, which is the future scope of this work. As our goal here is to examine the robustness of the learned representations by different pre-training techniques, we keep the encoder frozen. In summary, the combination of self-supervised pre-training with supervised fine-tuning offers a promising learning framework for fingerprint biometrics. Our methodology aims to leverage the strengths of both self-supervised and supervised learning, offering a pathway towards robust, efficient, and data-savvy fingerprint biometrics systems.

3 EXPERIMENTS

In this section, we discuss the experiments performed to evaluate our proposed approach's efficacy. We cover the specifics of our experimental setup, including the datasets used, the training details, and the evaluation metrics employed.

3.1 Datasets and Preprocessing

This study employs datasets comprising synthetic and real-world fingerprint images from SFinGe (Cappelli, 2004), FVC (Maio et al., 2002a; Maio et al., 2002b; Maio et al., 2004), and NIST SD-302 (Fiumara et al., 2019). SFinGe simulates real-world challenges, while FVC and NIST SD-302 offer large-scale, realistic fingerprint data for generalizability. This dataset combination enables model training and evaluation under diverse conditions. Synthetic data provides scalability and control, while real-world data ensures applicability. During self-supervised pre-training, only training dataset fingerprint images are used, without labels. Ground truth is needed for the enhancement task, obtained from clean images for SFinGe and generated for NIST SD-302 and FVC using a classical approach (Hong et al., 1998). The next phase involves a binary classification task for fingerprint verification. Data augmentation, vital for self-supervised learning, employs random transformations like rotation, color jitter, resize, crop, and Gaussian blur.

3.2 Implementation Details

We perform experiments using the PyTorch (Paszke et al., 2017) framework on an Nvidia GeForce RTX 2080 Ti GPU for training.

Our proposed enhancement-based pre-training utilizes the U-Net architecture. This U-Net encoder is employed consistently for pre-training with other self-supervised methods to ensure fair comparison. The U-Net has a depth of 5 layers, each with 2 convolutions, and expects 512 x 512-pixel grayscale fingerprint images. The encoder outputs a 4096-dimensional vector bottleneck, reduced to 512-d with an MLP projection head. Depth-wise convolutions minimize parameters. We use L_2 loss for U-Net's enhancement-based pre-training, adopting losses described in respective papers for other techniques.

For pre-training with existing self-supervised methods, a grid search identifies optimal hyperparameters. Models are pre-trained for 50 epochs with early stopping.

In probing experiments, MLP projection head weights are adapted for verification while keeping encoder weights fixed. We create 1:3 positive-to-negative pairs for training and testing verification sets from each dataset. After training, models are evaluated on test sets, reporting metrics like verification accuracy, precision, recall, and F1-score. Results are presented in two ways: 1) using the binary classifier over the MLP projection head (Figure 1-a) and 2) utilizing representations with thresholds on cosine sim-

Table 1: Enhancement pre-training stage results with U-Net architecture.

Method	SSIM	RSME	PSNR	NFIQ2
Raw Images	0.595	113.23	6.53	33.42
Enhancement U-Net	0.903	39.38	16.72	51.26

Figure 3: Degraded (top row) and Enhanced (bottom row) image pairs on FVC dataset from enhancement pre-training.

Table 2: Verification accuracy on SFinGe test dataset with genuine and imposter pairs.

SFinGe - Accuracy						
Method	Classification			Similarity		
	Imposter	Genuine	Entire Data	Imposter	Genuine	Entire Data
SimCLR	0.968	0.881	0.946	0.982	0.749	0.923
SimSiam	0.972	0.362	0.819	0.888	0.648	0.828
MoCo	0.963	0.881	0.942	0.955	0.845	0.927
BYOL	0.96	0.825	0.926	0.963	0.718	0.901
Ours	**0.982**	**0.886**	**0.958**	**0.975**	**0.847**	**0.943**

Table 3: F1 score on SFinGe test dataset with genuine and imposter pairs.

SFinGe - F1 score						
Method	Classification			Similarity		
	Imposter	Genuine	Entire Data	Imposter	Genuine	Entire Data
SimCLR	0.98	0.8	0.803	0.98	0.78	0.781
SimSiam	0.96	0.44	0.442	0.92	0.47	0.469
MoCo	0.98	0.79	0.785	0.97	0.74	0.737
BYOL	0.97	0.74	0.742	0.97	0.69	0.689
Ours	**0.99**	**0.86**	**0.858**	**0.98**	**0.81**	**0.821**

ilarity (Figure 1-b). The first method evaluates the model as an end-to-end verification network, while the second explores the potential of learned representations for similarity search and recognition tasks.

3.3 Results

The models are first pre-trained to learn fingerprint representations using the enhancement-based approach and various self-supervised learning strategies. Because these representations are not explicitly trained for fingerprint verification or identification, using them directly for evaluation is inappropriate. To gauge the stability and usefulness of these learned representations, we add linear layers to the frozen pre-trained encoders and then train the models for fingerprint verification tasks. The encoders remain frozen, allowing only the weights of the MLP to adjust to the task, keeping the original representations unchanged. This setup aids in comparing the efficacy of different self-supervised learning techniques against our method. The results of our probing experiments are presented in Table 4 (Verification Accuracy) and 5 (F1-score). The verification accuracy and F1-score on the SFinGe test set are shown in Tables 2 and 3 respectively. Figure 3 shows a few sample pairs of input and predicted images from the pre-trained U-Net model on the enhancement task used in our ap-

Table 4: Verification accuracy on FVC test dataset with genuine and imposter pairs.

FVC - Accuracy						
Method	**Classification**			**Similarity**		
	Imposter	**Genuine**	**Entire Data**	**Imposter**	**Genuine**	**Entire Data**
SimCLR	0.915	0.619	0.841	0.943	0.537	0.841
SimSiam	0.956	0.122	0.747	0.387	0.733	0.473
MoCo	0.902	0.522	0.807	0.896	0.56	0.812
BYOL	0.886	0.568	0.806	0.926	0.477	0.813
Ours	**0.957**	**0.73**	**0.900**	**0.933**	**0.818**	**0.904**

Table 5: F1 score on FVC test dataset with genuine and imposter pairs.

FVC - F1 score						
Method	**Classification**			**Similarity**		
	Imposter	**Genuine**	**Entire Data**	**Imposter**	**Genuine**	**Entire Data**
SimCLR	0.94	0.5	0.502	0.95	0.51	0.51
SimSiam	0.94	0.16	0.156	0.55	0.19	0.186
MoCo	0.93	0.42	0.417	0.92	0.43	0.431
BYOL	0.92	0.42	0.421	0.94	0.43	0.432
Ours	**0.97**	**0.68**	**0.679**	**0.96**	**0.66**	**0.659**

proach.

Our approach is compared with methods like SimCLR v2, SimSiam, MoCo v2, and BYOL on the SFinGe and FVC test sets for fingerprint verification. Verification accuracy serves as the evaluation metric for each method. The test data for fingerprint verification consists of a 1:3 ratio of positive to negative pairs, setting the random guess accuracy at 75%. Verification accuracy is measured in two ways as described before. This is presented in the below tables under the 'Classifier' column. The second way is represented under the 'Similarity' column in the tables. Moreover, we also report the ROC curves in Figure 4 for both datasets.

As seen from the results, our enhancement-based pre-training method consistently outperforms other self-supervised strategies across both test sets. SimCLRv2 also consistently performs well. SimSiam and BYOL methods show comparatively poor performance. It is noteworthy that all models perform better on the SFinGe test set than on the FVC test set. We believe this is due to two primary factors: the training sets contain more data from SFinGe than FVC, potentially resulting in a bias towards the former, and SFinGe is a synthetic dataset while FVC consists of real fingerprints, making the latter more challenging. Hence, the performance of models on FVC data is the real measure of the efficacy of models. Importantly, our method also provides superior performance when verification is based on the similarity of the representations, suggesting that the learned representations are also useful for fingerprint recognition.

4 LIMITATIONS AND FUTURE WORK

Despite promising results, our model demonstrates greater efficacy on the synthetic SFinGe dataset than on the real-world FVC dataset. This could be attributed to potential bias from underrepresentation of FVC data in training sets and complexities in real-life fingerprint data. Another limitation is the lack of specific training and evaluation for the fingerprint recognition task. While our model shows potential, a dedicated evaluation is essential for a comprehensive understanding of its performance. The effectiveness of self-supervised learning relies on data quality and diversity, and our study used linear probing, leaving room to explore alternative approaches like softmax or ArcFace-based classification.

Future work should address limitations by incorporating a more diverse set of real-world fingerprint datasets during training. Exploring the option of training the encoder with a smaller learning rate, rather than freezing it, could enhance generalizability and robustness. Specific training and evaluation for the recognition task, investigating alternative linear probing techniques, and exploring various self-supervised learning methods are valuable directions for further optimization.

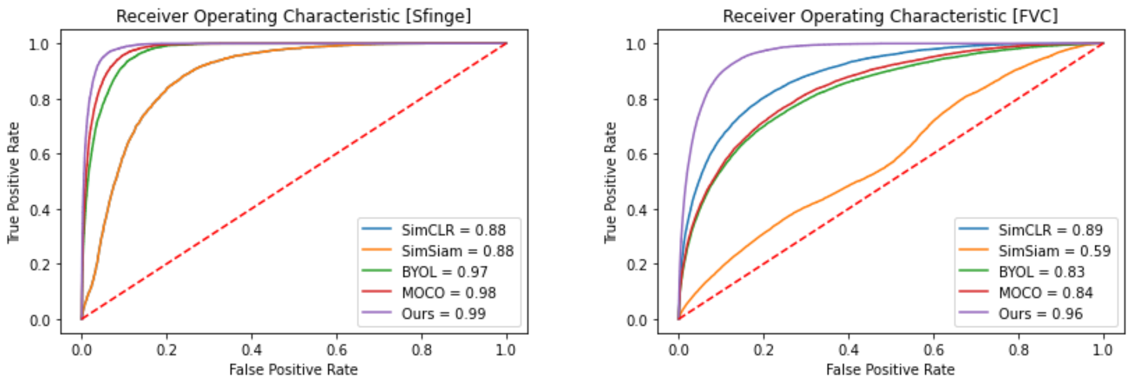

Figure 4: ROC curve based on similarity scores on SFinGe dataset(left) and FVC dataset (right).

5 CONCLUSION

In this study, we explored diverse self-supervised learning techniques to pre-train a model for effective fingerprint representations in recognition and verification. A novel approach involved leveraging fingerprint enhancement as a self-supervised pre-training method. Probing experiments assessed the effectiveness of learned representations across various pre-training strategies. Comparisons against SimCLR v2, SimSiam, MoCo v2, and BYOL methods on SFinGe and FVC datasets consistently showed our method's superior verification performance. Notably, our model excelled in similarity-based verification, underscoring its effectiveness in fingerprint recognition tasks. However, models performed better on the synthetic SFinGe dataset, hinting at potential bias in the training set and real-world data complexities. Future work will expand to diverse real-world fingerprint datasets, improving model generalizability. We'll also explore additional self-supervised methods for enhanced adaptability to real-world complexities, emphasizing the potential of self-supervised learning in fingerprint biometrics while pointing to areas for exploration and refinement.

REFERENCES

Allen, R., Sankar, P., and Prabhakar, S. (2005). *Fingerprint Identification Technology*, pages 22–61. Springer London, London.

Cappelli, R. (2004). Sfinge : an approach to synthetic fingerprint generation.

Cappelli, R., Ferrara, M., and Maltoni, D. (2010a). Minutia cylinder-code: A new representation and matching technique for fingerprint recognition. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2128–2141.

Cappelli, R., Ferrara, M., Maltoni, D., and Tistarelli, M. (2010b). Mcc: A baseline algorithm for fingerprint verification in fvc-ongoing. In *2010 11th International Conference on Control Automation Robotics & Vision*, pages 19–23. IEEE.

Cappelli, R., Maio, D., Lumini, A., and Maltoni, D. (2007). Fingerprint image reconstruction from standard templates. *IEEE transactions on pattern analysis and machine intelligence*, 29(9):1489–1503.

Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. (2021). Unsupervised learning of visual features by contrasting cluster assignments.

Chang, S.-H., Cheng, F.-H., Hsu, W.-H., and Wu, G.-Z. (1997). Fast algorithm for point pattern matching: invariant to translations, rotations and scale changes. *Pattern recognition*, 30(2):311–320.

Chen, T., Kornblith, S., Swersky, K., Norouzi, M., and Hinton, G. (2020a). Big self-supervised models are strong semi-supervised learners.

Chen, X., Fan, H., Girshick, R., and He, K. (2020b). Improved baselines with momentum contrastive learning.

Chikkerur, S., Govindaraju, V., and Cartwright, A. N. (2005). Fingerprint image enhancement using stft analysis. In *International Conference on Pattern Recognition and Image Analysis*, pages 20–29. Springer.

Darlow, L. N. and Rosman, B. (2017). Fingerprint minutiae extraction using deep learning. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE.

Deshpande, U. U., Malemath, V., Patil, S. M., and Chaugule, S. V. (2020). Cnnai: a convolution neural network-based latent fingerprint matching using the combination of nearest neighbor arrangement indexing. *Frontiers in Robotics and AI*, 7:113.

Engelsma, J. J., Cao, K., and Jain, A. K. (2019). Learning a fixed-length fingerprint representation. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):1981–1997.

Fiumara, G., Flanagan, P., Grantham, J., Ko, K., Marshall, K., Schwarz, M., Tabassi, E., Woodgate, B., and

Boehnen, C. (2019). Nist special database 302: Nail to nail fingerprint challenge.

Gavas, E. and Namboodiri, A. (2023). Finger-UNet: A u-net based multi-task architecture for deep fingerprint enhancement. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*. SCITEPRESS - Science and Technology Publications.

Greenberg, S., Aladjem, M., and Kogan, D. (2002). Fingerprint image enhancement using filtering techniques. *Real-Time Imaging*, 8(3):227–236.

Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al. (2020). Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284.

Gutmann, M. and Hyvärinen, A. (2010). Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings.

Hong, L., Wan, Y., and Jain, A. (1998). Fingerprint image enhancement: Algorithm and performance evaluation. *IEEE transactions on pattern analysis and machine intelligence*, 20(8):777–789.

Jain, A., Hong, L., and Bolle, R. (1997). On-line fingerprint verification. *IEEE transactions on pattern analysis and machine intelligence*, 19(4):302–314.

Jain, A., Ross, A., and Prabhakar, S. (2001). Fingerprint matching using minutiae and texture features. In *Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205)*, volume 3, pages 282–285. IEEE.

Jaiswal, A., Babu, A. R., Zadeh, M. Z., Banerjee, D., and Makedon, F. (2020). A survey on contrastive self-supervised learning. *Technologies*, 9(1):2.

Jing, L. and Tian, Y. (2020). Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4037–4058.

Kim, B.-G., Kim, H.-J., and Park, D.-J. (2002). New enhancement algorithm for fingerprint images. In *2002 International Conference on Pattern Recognition*, volume 3, pages 879–882. IEEE.

Liu, M., Chen, X., and Wang, X. (2014). Latent fingerprint enhancement via multi-scale patch based sparse representation. *IEEE Transactions on Information Forensics and Security*, 10(1):6–15.

Liu, M. and Qian, P. (2020). Automatic segmentation and enhancement of latent fingerprints using deep nested unets. *IEEE Transactions on Information Forensics and Security*, 16:1709–1719.

Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., and Tang, J. (2021). Self-supervised learning: Generative or contrastive. *IEEE transactions on knowledge and data engineering*, 35(1):857–876.

Maio, D., Maltoni, D., Cappelli, R., Wayman, J. L., and Jain, A. K. (2002a). Fvc2000: Fingerprint verification competition. *IEEE transactions on pattern analysis and machine intelligence*, 24(3):402–412.

Maio, D., Maltoni, D., Cappelli, R., Wayman, J. L., and Jain, A. K. (2002b). Fvc2002: Second fingerprint verification competition. In *2002 International Conference on Pattern Recognition*, volume 3. IEEE.

Maio, D., Maltoni, D., Cappelli, R., Wayman, J. L., and Jain, A. K. (2004). Fvc2004: Third fingerprint verification competition. In Zhang, D. and Jain, A. K., editors, *Biometric Authentication*, pages 1–7, Berlin, Heidelberg. Springer Berlin Heidelberg.

Maltoni, D., Maio, D., Jain, A. K., and Feng, J. (2009). *Fingerprint Matching*, pages 167–233. Springer London, London.

Maltoni, D., Maio, D., Jain, A. K., and Feng, J. (2022). *Fingerprint Sensing*, pages 63–114. Springer International Publishing, Cham.

Nguyen, D.-L., Cao, K., and Jain, A. K. (2018). Robust minutiae extractor: Integrating deep networks and fingerprint domain knowledge. In *2018 International Conference on Biometrics (ICB)*, pages 9–16. IEEE.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.

Qian, P., Li, A., and Liu, M. (2019). Latent fingerprint enhancement based on denseunet. In *2019 International Conference on Biometrics (ICB)*, pages 1–6. IEEE.

Rahman, S. M., Ahmad, M. O., and Swamy, M. (2008). Improved image restoration using wavelet-based denoising and fourier-based deconvolution. In *2008 51st Midwest Symposium on Circuits and Systems*, pages 249–252. IEEE.

Ratha, N. K., Karu, K., Chen, S., and Jain, A. K. (1996). A real-time matching system for large fingerprint databases. *IEEE transactions on pattern analysis and machine intelligence*, 18(8):799–813.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks.

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.

Sherlock, B., Monro, D., and Millard, K. (1992). Algorithm for enhancing fingerprint images. *Electronics letters*, 18(28):1720–1721.

Tang, Y., Gao, F., Feng, J., and Liu, Y. (2017). Fingernet: An unified deep network for fingerprint minutiae extraction. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 108–116. IEEE.

Wayman, J., Jain, A., Maltoni, D., and Maio, D. (2005). *An Introduction to Biometric Authentication Systems*, pages 1–20. Springer London, London.

Yang, J., Liu, L., and Jiang, T. (2002). Improved method for extraction of fingerprint features. In *Second International Conference on Image and Graphics*, volume 4875, pages 552–558. SPIE.

Zaeri, N. (2011). Minutiae-based fingerprint extraction and recognition. *Biometrics*.

SCIENCE AND TECHNOLOGY PUBLICATIONS