

# TattTRN: Template Reconstruction Network for Tattoo Retrieval

Lazaro Janier Gonzalez-Soler<sup>1</sup>, Maciej Salwowski<sup>2</sup>, Christian Rathgeb<sup>1</sup> and Daniel Fischer<sup>1</sup>,  
<sup>1</sup>da/sec - Biometrics and Security Research Group, Darmstadt, Germany

{lazaro-janier.gonzalez-soler, christian.rathgeb, daniel.fischer}@h-da.de

<sup>2</sup> Faculty of Computer Science, Technical University of Denmark, Denmark

s223525@student.dtu.dk

## Abstract

*Tattoos have been used effectively as soft biometrics to assist law enforcement in the identification of offenders and victims, as they contain discriminative information, and are a useful indicator to locate members of a criminal gang or organisation. Due to various privacy issues in the acquisition of images containing tattoos, only a limited number of databases exists. This lack of databases has delayed the development of new methods to effectively retrieve a potential suspect's tattoo images from a candidate gallery. To mitigate this issue, in our work, we use an unsupervised generative approach to create a balanced database consisting of 28,550 semi-synthetic images with tattooed subjects from 571 tattoo categories. Further, we introduce a novel Tattoo Template Reconstruction Network (TattTRN), which learns to map the input tattoo sample to its respective tattoo template to enhance the distinguishing attributes of the final feature embedding. Experimental results with real data, i.e. WebTattoo and BIVTatt databases, demonstrate the soundness of the presented approach: an accuracy of up to 99% is achieved for checking at most the first 20 entries of the candidate list.<sup>1</sup>*

## 1. Introduction

The use of tattoos as soft biometrics to assist law enforcement in identifying suspects has steadily grown along with their popularity in society. In 2015, The National Institute of Standards and Technology (NIST) (Tatt-C) [31] reported that one-fifth of US adults have at least one tattoo, making the US population the third most tattooed in the world, after Italy and Sweden. This trend was constantly expanding, as evidenced by a survey conducted in 2021<sup>2</sup>, which revealed that 26% of Americans have at least one tattoo. In contrast

<sup>1</sup><https://github.com/ljsoler/TattTRN>

<sup>2</sup><https://www.statista.com/statistics/259598/share-of-americans-with-at-least-one-tattoo/>

Table 1. Overview of available tattoo databases.

Database	#Categories	#Samples	Public	Semi-synthetic
HDA-STD [8]	-	5,500	Yes	✓
WebTattoo [11]	400*	1,400	Yes	✗
DeMSI [14]	-	890	Yes	✗
Tatt-C [31]	157†	215	No	✗
BIVTatt [33]	210	4,200	Yes	✗
PinTatt [33]	160	454	No	✗
NTU-Tattoo-V1 [45]	-	10,000	Yes	✗
<b>Ours</b>	<b>571</b>	<b>28,550</b>	<b>Yes</b>	<b>✓</b>

\* Publicly available training set.

† Number of categories reported for the identification case.

to biometric characteristics, e.g. fingerprints and faces, tattoos cannot be used to directly establish the identity of a subject. However, tattoos, unlike other soft biometrics such as gender, age or race, contain more discriminative information to support the identification of individuals and are a useful indicator to track members of a criminal gang or organisation [31]. Therefore, tattoo recognition represents an area of interest for forensic investigators, which motivates the development of automated image-based tattoo retrieval techniques [30].

In the context of tattoo retrieval, the NIST Tatt-C and Tatt-E challenges advanced the development of tattoo detection and identification systems for real-world application scenarios [32]. Earlier tattoo image retrieval practices were based on keyword or metadata matching, where law enforcement agencies typically followed the ANSI/NIST-ITL 1-2000 standard and assigned a single keyword to each tattoo image in the database [11]. The limitations of keyword-based systems (e.g., limited and insufficient vocabulary to describe different tattoo patterns and inconsistent labels) led to the development of techniques that represented tattoos as handcrafted descriptors capturing the texture, appearance, or colour of the tattoos [10, 24, 26].

With the introduction and success of deep neural networks (DNNs) in various pattern recognition and computer vision applications, some recent research efforts have been devoted to tattoo localisation [41], detection [3, 41] or seg-



Figure 1. Low-quality tattoos.

Figure 2. Examples of tattoos from WebTattoo [11] representing a challenge for state-of-the-art solutions.

mentation [8]. Moreover, a few techniques have focused on tattoo retrieval or identification through the representation of tattoos as compact binary [11] or floating-point embeddings [6, 7, 34]. The main reasons that have slowed down the development of new tattoo retrieval methods are, on the one hand, the lack of large-scale tattoo databases, as shown in Tab. 1. It should be noted that a few tattoo databases are publicly available and that most of them were designed for tattoo detection [45] and segmentation [8, 14] or provide a limited number of subjects [33] or samples [11]. On the other hand, previous works often neglect challenges that cause performance degradation, *e.g.* low-image quality, see Fig. 2. Consequently, a moderate identification rate (IR) below 65% at the rank-1 is reported in the latest published tattoo identification pipelines.

To overcome the aforementioned challenges, we introduce a large-scale semi-synthetic tattoo database based on which a novel tattoo retrieval approach is proposed that transforms tattooed human skin into the respective tattoo template to improve the final tattoo description. Thereby, the accuracy of tattoo retrieval can be significantly improved. The main contributions of this scientific work are:

- A large-scale semi-synthetic tattoo database that consists of 28,550 samples from 571 different tattoo templates or categories. To simulate a real-life scenarios, tattooed samples are generated and augmented in terms of scale, colour adjustments, distortions, and opacity. For the synthesis, the general requirements for generating synthetic biometric data are considered [28].
- A tattoo retrieval solution that transforms tattooed human skin into the corresponding tattoo template. The reconstruction function optimised together with a well-known identity loss function alleviates the above issues encountered in real images and, thus, improves the identification performance of state-of-the-art algorithms.
- A comprehensive evaluation of the proposed system on real images according to the metrics defined in the international standard ISO/IEC 19795-1 [19] for biometric testing and reporting. Experimental results conducted on WebTattoo [11] and BIVTatt [33] show a convincing performance improvement of the proposed system over the baselines.

The remainder of this paper is organised as follows: a brief review of existing tattoo retrieval systems is provided in Sect. 2. In Sect. 3, the fundamentals of the proposed approach are presented, along with a description of the tattoo generation. The experimental setup is explained in Sect. 4 and results as well as derived findings are discussed in Sect. 5. Finally, conclusions and future work directions are presented in Sect. 6.

## 2. Related Work

Among the various soft biometric traits, tattoos have received considerable attention by forensic investigators in recent years, due to their prevalence among the criminal sector of the population and their prominence in visual attention. For more than 5,000 years, humans have marked their bodies with tattoos to express personal beliefs or to associate themselves with a group. The oldest evidence of tattooing was found on the body of Ötzi, the Iceman. Ötzi, Europe’s most famous mummy, was discovered by German hikers in the Alps in 1991. The next oldest evidence of tattooing comes from mummies believed to have died between 3351 and 3017 BC in Ancient Egypt<sup>3</sup>. In forensics, tattoos have also proved useful in helping to identify victims of terrorist attacks such as 9/11 and natural disasters such as the 2004 Indian Ocean tsunami [20].

To assist in the identification of subjects, early tattoo retrieval systems relied on keyword or metadata matching. For this purpose, law enforcement agencies usually followed the ANSI/NIST-ITL 1-2000 standard and assigned a single keyword to each tattoo image in the database [29]. Since these solutions had obvious drawbacks, for example, *i)* a limited vocabulary to describe numerous tattoo patterns, *ii)* several keywords could be used to adequately describe a tattoo, and *iii)* inconsistency in the labelling of tattoos, the next set of retrieval approaches focused on the use of hand-crafted descriptors to represent the colour, texture, shape, and appearance of tattoos [10, 24, 26]. A comprehensive review of these methods up to 2019 can be found in [11].

In this work, we restrict to briefly summarising deep learning-based tattoo retrieval solutions due to their reported success and outstanding performance in various pattern recognition and computer vision tasks [21]. In 2019, Han *et al.* [11] introduced a system that was able to learn tattoo detection and compact tattoo representation jointly in a single DNN by multitask learning. Following this idea, Zhang *et al.* [47] proposed a DNN-based framework for joint tattoo detection and re-identification of individuals. Their reported IRs improved in the WebTattoo database [11] in case the detection module was activated. Nicolás-Díaz *et al.* [33] evaluated publicly available DNN models, pre-

<sup>3</sup><https://www.nationalgeographic.com/history/article/tattoos-mummies-ancient-cultures-symbols-meaning>

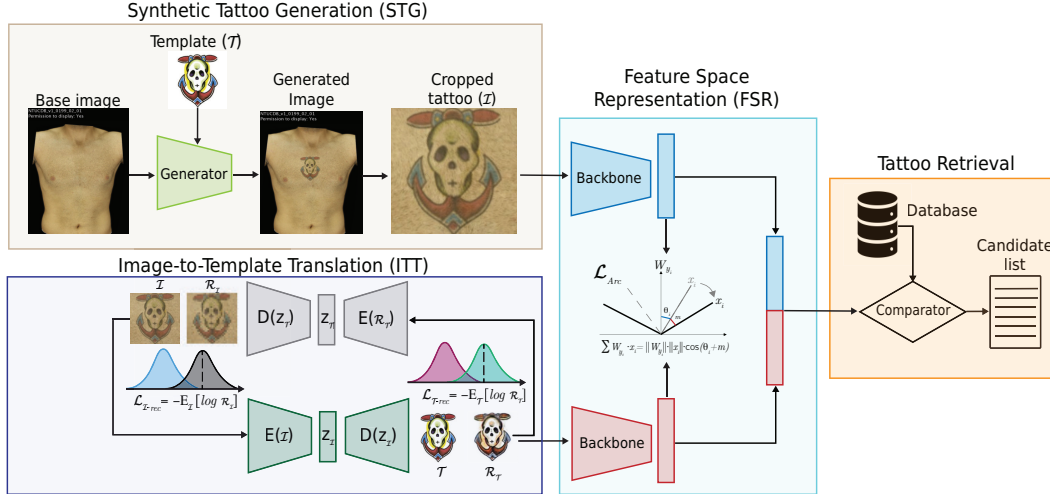


Figure 3. Conceptual overview of the proposed system: the template transformed from the input image helps to mitigate challenges related to the capturing process. The final feature embedding representing the salient properties of the input tattoo is the concatenation of the two computed feature embeddings and can be used to retrieve similar tattoo samples in a database.

trained with large generic image databases, for tattoo identification and showed that these DNNs can achieve high IRs without even fine-tuning them for the target task of tattoo retrieval. The same authors also made the BIVTatt database publicly available in [33]. In addition, Nicolás-Díaz *et al.* [34] presented an attention pooling mechanism to adapt intermediate convolutional layers of pre-trained DNNs for the tattoo identification task.

The aforementioned tattoo retrieval systems are based on deep learning techniques, which require a considerable amount of training data to achieve good performance. In Tab. 1, we summarise the main features of the databases used for such purposes. Note that, on the one hand, that Tatt-C and PinTatt are not available to the research community. On the other hand, the existing databases only consist of a few tattoo categories or samples which do not cover most real-life scenarios. This could make algorithms trained on those databases prone to overfitting.

Image synthesis has recently emerged as a reliable solution to both privacy concerns and the lack of training data [21]. To our knowledge, few studies have addressed the (semi-)synthetic generation of tattoo images for retrieval purposes. Existing methods to create face images with tattoos [18], are mainly designed for the visualisation of tattoos in controlled skin images [2], or tattoo segmentation [8].

### 3. Tattoo Template Reconstruction Network

To mitigate the aforementioned problems, we introduce our Tattoo Template Reconstruction Network (TattTRN), which learns to map the input tattoo sample to its respective tattoo template. This tattoo template mapping is then used

to enhance the distinctive attributes of the final feature embedding. Fig. 3 shows the conceptual overview of the proposed TattTRN. In our work, we hypothesise that a clean tattoo template may contain meaningful information to assist both an supervised learning approach and forensic investigators in making decisions. In the case of low-quality images, such as those shown in Fig. 2, where the tattoo is almost imperceptible to the human eye, the proposed concept could reconstruct the corresponding template and, thus, help forensic investigators. In addition to the above benefits, the transformation of an input tattoo image into a clean template would also provide subject identity suppression, as the reconstructed template is not expected not contain traces of skin colour or other sensible information.

The following subsections describe the different components of our proposed TattTRN (Sect. 3.1), as well as the definition and synthesis of the tattoo database (Sect. 3.2) involved in training the approach.

#### 3.1. Main Components

The proposed approach consists of three main components (see Fig. 3): synthetic tattoo generation (STG), image-to-template translation (ITT) and feature space representation (FSR), which can be used for tattoo retrieval. STG allows the offline/online generation of tattooed samples, which can be directly represented as a feature embedding or translated into the predefined tattoo template. Whereas a direct representation of the sample (light blue trapezoid in FSR) aims at capturing intrinsic properties of the input image ( $\mathcal{I} \in \mathbb{R}^{n \times n}$ ), the ITT module focuses on the construction of the respective clean tattoo template ( $\mathcal{R}_{\mathcal{T}}$ ), which is, in turn, represented as a feature embedding (light red trapezoid

in FSR). Both feature embeddings contain prominent characteristics of the input image and complement each other to enrich the final representation of the tattoos. For tattoo retrieval (a use case represented by the yellow box on the right in Fig. 3), we use the concatenation of both embeddings to form a feature vector of size  $2 \cdot K$ . To translate the input image into a clean tattoo template in the ITT component, we use, as a proof-of-concept, the Unet [40] network based on the ResNet34 encoder [12] and built a cyclic translation, as done in CycleGAN [48]. Thus, the quality of the template reconstruction is improved. Note that other encoder-decoder architectures proposed for image-to-image translation, such as Diffusion Models [25], could be employed, and thus improve the reconstruction results yielded by Unet. To optimise the cyclic Unet, two binary cross-entropy (BCE) loss functions are computed and combined as follows:

$$\begin{aligned} \mathcal{L}_{\mathcal{T}-rec} &= -\mathbf{E}_{\mathcal{T}}[\log \mathcal{R}_{\mathcal{T}}] \\ &= -[\mathcal{T} \cdot \log \mathcal{R}_{\mathcal{T}} + (1 - \mathcal{T}) \cdot \log(1 - \mathcal{R}_{\mathcal{T}})], \end{aligned} \quad (1)$$

$$\begin{aligned} \mathcal{L}_{\mathcal{I}-rec} &= -\mathbf{E}_{\mathcal{I}}[\log \mathcal{R}_{\mathcal{I}}] \\ &= -[\mathcal{T} \cdot \log \mathcal{R}_{\mathcal{I}} + (1 - \mathcal{T}) \cdot \log(1 - \mathcal{R}_{\mathcal{I}})], \end{aligned} \quad (2)$$

$$\mathcal{L}_{rec} = \mathcal{L}_{\mathcal{T}-rec} + \mathcal{L}_{\mathcal{I}-rec} \quad (3)$$

where  $\mathcal{I}$  is the semi-synthetic generated image using the tattoo template  $\mathcal{T}$  in STG and  $\mathcal{R}_{\mathcal{I}}$  is the reconstructed image from the translated tattoo template  $\mathcal{R}_{\mathcal{T}}$  in ITT. BCE measures the difference between two probability distributions, *i.e.*,  $(\mathcal{T}, \mathcal{R}_{\mathcal{T}})$  and  $(\mathcal{I}, \mathcal{R}_{\mathcal{I}})$  in this case and shows advantages in terms of convergence for image reconstruction over other loss functions such as the Minimum Squared Error (MSE).

On top of the encoder-decoder network, a backbone network computes the embedding representation (light red trapezoid in FSR, Fig. 3) of the reconstructed template  $\mathcal{R}_{\mathcal{T}}$ . In our work, both backbones (light blue and red trapezoids of FSR) are based on the same architecture and compute a feature embedding of size  $K$  each. However, they do not share weights, as they process different representations of the input tattoo. Whereas the upper backbone (light blue trapezoid) of FSR is fed by the semi-synthetic raw tattoo image generated by the STG component, the lower backbone (light red trapezoid) computes the feature embedding of the clean tattoo template.

Together with  $\mathcal{L}_{rec}$  we select the ArcFace loss [5] to optimise the backbones separately. Angular margin penalty-based softmax loss was initially proposed to optimise face recognition systems [1, 5]. It aims to extend the softmax decision boundary to improve intra- and inter-class variation by applying an angular penalty margin on the angle between

feature embeddings and their respective weights. The ArcFace loss has reported high performance in the main benchmarks and is defined as follows [5]:

$$\mathcal{L}_{Arc} = -\frac{1}{N} \sum_{i \in N} \log \frac{e^{s \cdot \cos(\theta_{y_i} + m)}}{e^{s \cdot \cos(\theta_{y_i} + m)} + \sum_{j=1, j \neq y_i}^C e^{s \cdot \cos \theta_j}}, \quad (4)$$

where  $N$  and  $C$  represent the training batch size and the number of classes (*i.e.*, tattoo categories in our article), respectively,  $y_i \in \{1, C\} \in \mathbb{N}$  is the tattoo category of the  $i^{th}$  sample and  $\theta_{y_i}$  is the angle between the tattoo feature embedding  $x_i$  and the respective weight  $w_{y_i}$ . Note that linear function  $x_i \cdot w_{y_i}$ <sup>4</sup> can be represented as  $\|x_i\| \cdot \|w_{y_i}\| \cdot \cos \theta_{y_i}$ . In our work, both  $x_i$  and  $w_{y_i}$  are normalised and hence  $\|x_i\| = 1$  and  $\|w_{y_i}\| = 1$ , resulting only in the softmax optimisation of  $\cos \theta_{y_i}$  along with the angular margin  $m$  *i.e.*,  $\cos(\theta_{y_i} + m)$ .  $s$  is the scale factor according to [44].

Finally, the loss combining both objective functions (*i.e.*, Eq. 3 and Eq. 4) to optimise our proposed TattTRN approach is defined as:

$$\mathcal{L} = \frac{\mathcal{L}_{\mathcal{I}-Arc} + \mathcal{L}_{\mathcal{T}-Arc} + \lambda \cdot \mathcal{L}_{rec}}{3}, \quad (5)$$

where the hyperparameter  $\lambda$  was empirically set to 4 and aims to keep the independently produced losses from both objective functions in equilibrium.  $\mathcal{L}_{\mathcal{I}-Arc}$  and  $\mathcal{L}_{\mathcal{T}-Arc}$  represent the ArcFace losses for the backbones optimised over  $\mathcal{I}$  and  $\mathcal{R}_{\mathcal{T}}$ , respectively.

### 3.2. Tattoo Generation

Synthetic image generation has mainly focused on facial characteristics. In terms of market, face recognition has maintained a stable growth and its use has spread to many application contexts, such as financial transactions, border control and video surveillance. To synthesise high-quality facial images, several approaches based on GAN or diffusion models, such as the StyleGAN family [22, 23] and latent diffusion models [39], have been proposed in recent years. In comparison to the face, tattoo synthesis has received little attention. Recently, Andrej Karpathy demonstrated the robustness of diffusion models for generating realistic tattooed human images<sup>5</sup>. Ibsen *et al.* [17] proposed a tattoo generator to blend predefined tattoo templates with real faces and evaluated the performance impact of facial recognition systems on tattooed faces. Following this idea, Gonzalez-Soler *et al.* [8, 9] extended the previous pipeline to mix the predefined tattoo templates with any area of human skin and demonstrated its utility for the segmentation

<sup>4</sup>The bias parameter of the fully connected layer representing the feature embedding is set to 0

<sup>5</sup><https://www.youtube.com/watch?v=sM9bozW295Q>

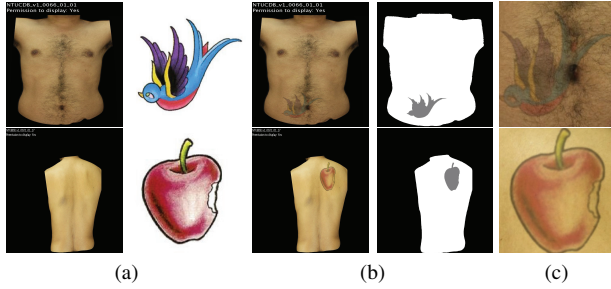


Figure 4. Examples of tattoos generated on chest and back images with their respective segmentation maps (4b) using base images and random tattoo templates (4a). Cropped tattoo images used in the network training (4c).

of real tattoos. In this article, we reuse the ideas shown in [8, 9] to generate 28,550 images of 571 different tattoo templates, *i.e.*, 50 images per template. The NTU-Back [36, 37] and NTU-Chest [35, 37] databases, consisting of 647 and 434 images respectively, are used as base images of human skin to blend with the tattoos. Randomly 25:25 images from both databases are used as base images to synthesise the 50 images per tattoo template. Fig. 4 shows examples of tattooed human skins. To simulate real images, the tattoo is made more realistic by adjusting the colour, and Gaussian blur, and reducing the opacity. Since the segmentation map is available, the final tattooed samples are then cropped to train the TattTRN system.

#### 4. Experimental Setup

The experimental evaluation goals are manifold: *i)* study the utility of the proposed semi-synthetic database for training a tattoo retrieval system capable of retrieving real tattoo samples, *ii)* establish a benchmark of the proposed TattTRN systems with baseline approaches, and *iii)* evaluate the capability of the TTE subsystem to enhance low-quality images. In all experiments, we follow a cross-database protocol where the proposed semi-synthetic database is used to train the systems and the real ones for evaluating performance. The test set is randomly divided into five subsets of biometric enrolment and identification transactions following a closed-set scenario, *i.e.*, searched identities are always in the enrolment database. Therefore, the mean identification rates alongside the standard deviation are presented as Cumulative Matching Characteristic (CMC) curves. Formally, CMC is a graphical presentation of the results of mated searches in a closed-set identification test, which plots the true positive identification rate (IR) as a function of a rank value [19]. The identification performance of the systems is also evaluated for an open set scenario in terms of False Negative Identification Rates (FNIR) and False Positive Identification Rates (FPIR), and their values are presented as DET curves [19].



Figure 5. Examples of the database used to evaluate the proposed TattTRN system.

#### 4.1. Databases

To evaluate the utility of both the semi-synthetic database and the proposed system, we selected the two publicly available databases in Tab. 1 that provide tattoo category labels *i.e.*, WebTattoo [11] and BIVTatt [33].

According to [11], WebTattoo comprises around 300K samples of 600 different tattoo categories which were extracted from the internet. For training, the authors randomly selected about 1,400 tattoo images from 400 tattoo classes and the remaining tattoo images from 200 tattoo classes were used for testing. In our experiments, we only used the training set that was made available to the research community.

BIVTatt [33] contains 210 original tattoo images and 4,200 images generated after applying 20 different types of transformations to the original images. Along with the images, the authors provided the bounding box information for each tattoo. Given the coordinate problems associated with the bounding boxes, only correctly cropped samples were selected, resulting in 3,103 samples from 159 tattoo categories. Fig. 5 shows examples of tattoo images in the WebTattoo [11] and BIVTatt [33] databases.

#### 4.2. Implementation Details

Both the proposed TattTRN systems and baselines were implemented in PyTorch [38] and trained utilising a Nvidia A100 Tensor Core GPU with 40 GB of GPU Memory over the generated 28,550 semi-synthetic images. The image size was set to  $224 \times 224$ . To further cover the feature space of real tattoo images, the brightness, contrast, saturation, and hue of the training images are randomly adjusted. In addition, the networks were initialised with their pre-trained weights on ImageNet [4] and trained for 100 epochs using the Adam optimiser with a learning rate of  $1e-5$  and weight decay of 0.95. A batch size of 64 images is also set for training. As backbones, the large ver-

Table 2. Closed-set identification performance of the proposed TattTRN approach for different backbones in terms of Rank-1 (%) on real images in WebTattoo [11] and BiVTatt [33]. The best result per backbone and database is highlighted in bold.

Database	$m$	WebTattoo									BiVTatt								
		0.1			0.5			0.9			0.1			0.5			0.9		
Backbone/ $K$		128	256	512	128	256	512	128	256	512	128	256	512	128	256	512	128	256	512
MobileNetv3		70.98	73.27	76.17	73.27	75.58	<b>76.51</b>	66.39	71.89	63.73	94.40	93.58	94.40	94.21	94.40	<b>94.65</b>	91.95	92.33	89.25
ResNet101		70.47	73.84	78.13	74.25	76.63	<b>78.45</b>	60.52	63.17	63.46	93.46	94.09	94.59	93.65	<b>94.78</b>	94.47	88.43	88.62	87.92
DenseNet121		73.59	<b>78.67</b>	76.78	74.64	77.27	78.13	75.01	77.03	76.03	93.14	95.16	93.96	93.33	94.28	94.72	94.72	94.53	<b>95.22</b>
EfficientNetv2		79.43	78.92	<b>80.44</b>	79.19	79.61	79.63	71.99	59.12	62.04	94.47	95.41	96.29	95.41	<b>96.54</b>	96.16	93.08	82.58	86.73
Swin		77.76	78.80	<b>81.60</b>	77.84	78.82	80.57	69.26	66.51	69.88	95.35	94.47	<b>96.54</b>	93.96	94.47	95.91	89.37	85.91	85.41
Avg.		74.45	76.70	78.62	75.84	77.58	78.66	68.63	67.54	67.03	94.16	94.54	95.16	94.11	94.89	95.18	91.51	88.79	88.91

sion of MobileNetv3 [13], 101-layer ResNet [12], 121-layer DenseNet [16], and the small version of EfficientNetv2 [42] and SwinTransformer (Swin) [27] are used. These networks offer coverage of the main approaches proposed in the last decade *e.g.*, attention mechanisms [15] and vision transformers [43]. Note that other architectures developed specifically for tattoo retrieval *e.g.*, the joint detection and compact representation scheme [11] or the weighted average pooling-based approach [34]<sup>6</sup>, can also be applied.

## 5. Results and Discussion

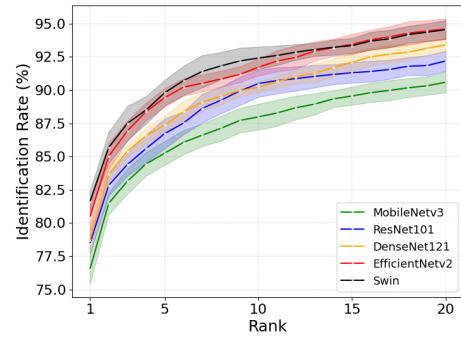
The result’s discussion relies on the above three goals: the closed-set identification performance of the proposed TattTRN approach (Sect. 5.1), as well as the respective open-set identification results 5.2 are reported. Sect. 5.3 shows a comparison of TattTRN with the baselines. To evaluate the utility of TattTRN, the benchmark is established against the direct embedding representation of the semi-synthetic tattoo images generated by the STG component. Finally, some examples of translating real images into templates are presented in Sect. 5.4.

### 5.1. Closed-set Evaluation


Since the use of ArcFace loss depends on the optimisation of the angular margin parameter (*i.e.*,  $m$ ), we report the mean closed-set identification performance of TattTRN for three values of  $m = \{0.1, 0.5, 0.9\}$  and three embedding sizes  $K = \{128, 256, 512\}$  in Tab. 2.  $K$  values greater than 512 would result in a final feature vector of  $2 \cdot K > 1024$ , leading to an efficiency deterioration of the system.

Note that the generated semi-synthetic tattoo images reflect the main properties of the real tattooed samples and their use allows our TattTRN approach to achieve high performance for the evaluated real databases. TattTRN is able to register an average IR of at most 81.60% at rank-1 for the challenging WebTattoo dataset and 96.54% at the same rank value for BiVTatt. Compared to the results reported in their respective articles, TattTRN reports a performance improvement of approximately 18% for Web-

<sup>6</sup>The implementation of both approaches has not been made publicly available.



(a) WebTattoo



(b) BiVTatt

Figure 6. CMC curves for different backbones combined with TattTRN.

Tattoo (*i.e.*, IR  $\approx$  64%) [11] and 25% for BiVTatt (*i.e.*, IR = 70.91%) [33]. It is worth mentioning that the benchmarking systems referred to in the articles of both databases were fully trained on real tattoo samples following an internal database protocol, in contrast to TattTRN, representing the most salient properties of the semi-synthetic images in the training to identify real specimens. In terms of parameter optimisation, we observe that most backbones achieve on average the best performance for  $K = 512$  and  $m = 0.5$  or  $m = 0.1$ , with the combination of TattTRN with Swin [27] being the best performing approach for WebTattoo (*i.e.*,

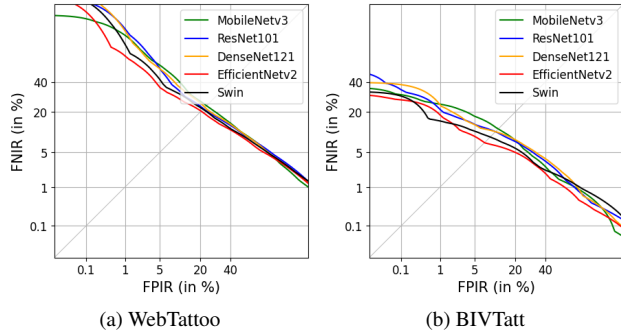


Figure 7. DET curves for different backbones combined with TattTRN.

IR = 81.60%) and BIVTatt (*i.e.*, IR = 96.54%) in rank-1.

Since in forensic investigations not only the first positions in the candidate lists are important, we also show the CMC curves for the different backbones combined with the TattTRN approach proposed in Fig. 6 using the best parameter configurations (*i.e.*, bold values in Tab. 2). Note that TattTRN combined with Swin achieves only the best IR in rank-1 for both databases. For rank values above 15, we can see that the EfficientNetv2 backbone yields an average IR of about 95% for WebTattoo and above 99% for BIVTatt. This implies that forensic investigators can detect a perpetrator based on tattoos with an accuracy of up to 99% by checking the first 20 entries of the candidate list.

## 5.2. Open-set Evaluation

We report the TattTRN performance in Fig. 7 for the scenario where the tattoo template or category is probably not included in the enrolment set (*i.e.*, open-set scenario). Note that TattTRN combined with different backbones yields similar equal error rates (EER) for WebTattoo. However, in line with the closed-set results shown in Sect. 5.1, EfficientNetv2 achieves the best identification performance for high-security thresholds for both databases. In particular, this backbone achieves an FNIR of more than 70% for WebTattoo and an FNIR of approximately 30% for BIVTatt for an FPIR = 0.1%. The latter result implies that, at most, 30 out of 100 mated transactions are not included in the candidate list if the system accepts only 1 out of 1000 non-mated transactions in the candidate list.

## 5.3. Benchmark of TattTRN

We evaluate the benefits of the translation module (*i.e.* ITT) for TattTRN. For this purpose, the entire TattTRN architecture is compared with the submodule comprised of only the raw embedding representation of the generated semi-synthetic tattoo image (*i.e.* light blue trapezoid in Fig. 3). Fig. 8 shows the benchmark of the whole TattTRN against the raw embedding computed by EfficientNetv2 [42]. Note




Figure 8. Benchmark of TattTRN against the respective best-performing backbone.

that the TattTRN benefits from the ITT component, resulting in a performance improvement for EfficientNetv2 on both databases. Specifically, for the challenging WebTattoo dataset, TattTRN outperforms EfficientNetv2 for all rank values and achieves an IR approximate of 95% in rank-20 *i.e.*, the perpetrator can be identified with an accuracy of up to 95% by checking at most the first 20 entries of the candidate list. Despite the advantages of the ITT component, there are still several images in WebTattoo whose quality should be further improved (see Sect. 5.4). It should be mentioned that extreme sharpening or blurring transformations to decrease image quality, as well as approaches that cut out part of the image (*e.g.*, CutMix [46]), were not taken into account in the training. The use of these operations could also increase the reconstruction performance of IIT, and hence of TattTRN.

## 5.4. Tattoo Retrieval Examples

Finally, in Fig. 9 we show some examples of candidate lists retrieved by TattTRN from searched images and the respective translated templates. Note that TattTRN is able to correctly build the template from those images that were successfully retrieved in the first positions of the candidate list (*i.e.*, images in Fig. 9a). The ITT component is therefore capable of encoding challenging patterns, such as the woman and the rose in the last two rows of Fig. 9a. It should be noticed that the remaining retrieved images, ranked from 2 to 4 in rows 2 and 3 of Fig. 9a, do not belong to the consulted tattoo category: TattTRN reports a similarity below 0.35.

Fig. 9b, on the other hand, shows some images for which

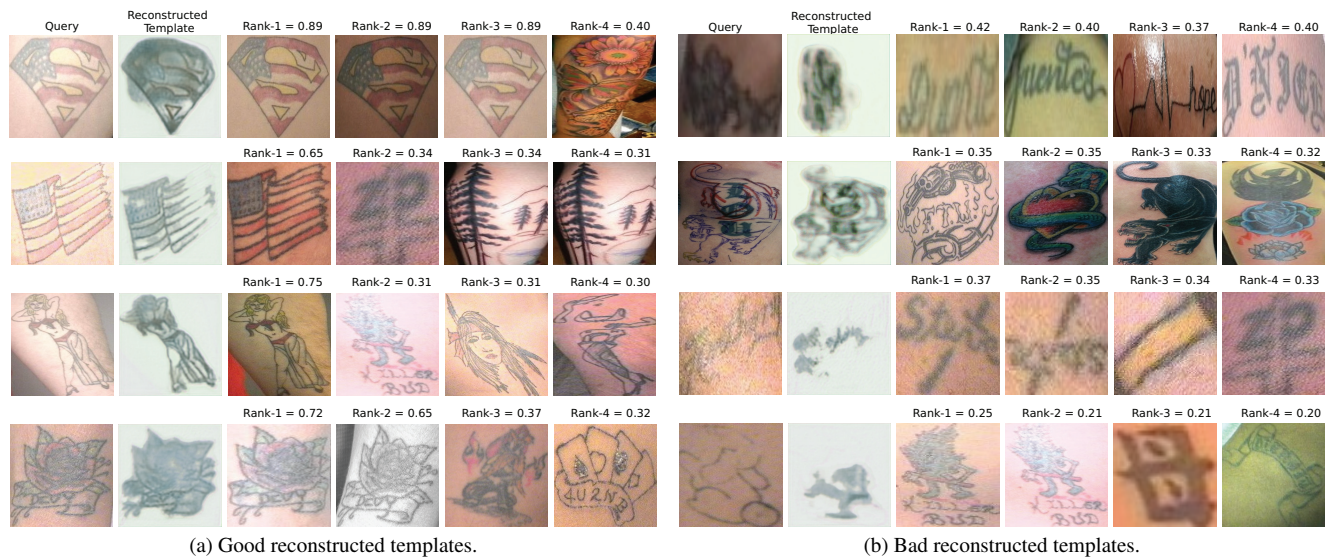


Figure 9. Example of correctly and wrongly retrieved tattooed samples from WebTattoo [11] for rank values between 1 and 4. The cosine similarity values are given for each case.

TattTRN could not correctly construct their clean tattoo template, resulting in lower similarity values for the first positions in the candidate list. Note that they either have extremely low image quality (*e.g.*, images of rows 1 and 3), which might even be difficult for human analysis, or they contain overlapping tattoos (*e.g.*, image of the second row). In the latter case, artistic characters are displayed in the background and an in-process panther in the foreground: TattTRN was able to partially construct its clean tattoo template. The overlapping of tattoos resulted in TattTRN regaining first place in the ranking for a tattoo containing Latin characters and third place for the corresponding correct tattoo category (black panther). Notice that all the tattoos retrieved in Fig. 9b have a similarity of less than 0.42 with respect to the queried tattoo samples. This implies, in line with the results in Fig. 7a, that the input image would be a false negative depending on the system threshold (typically 0.5) and therefore all items in the candidate list would be rejected by the system, even if there are items of the same category as the input image at the former positions in the candidate list (*e.g.*, the black panther). A potential solution to improve the results of TattTRN on the above images would be based on the use of a prompt-guided approach that also encodes the semantic meaning of the tattoo and the overlap between tattoos, or the combination with a super-resolution method that further improves the image quality of the input tattoos.

## 6. Conclusions

This work proposes a semi-synthetic tattoo database in conjunction with a tattoo retrieval approach called TattTRN,

which exploits the transformation of the input image into a clean tattoo template to enrich the final feature embedding and thus improve the retrieval performance of difficult images in two freely available databases: WebTattoo and BIVTatt. The experimental evaluation of TattTRN showed that the proposed balanced database, consisting of 28,550 images from 571 tattoo categories, gathers the main properties of the real images, resulting in an IR of up to 99% in the top 20 positions of the candidate list. Experimental results also reported high performance when TattTRN is combined with EfficientNetv2, yielding an IR of approximately 81% for WebTattoo in rank-1. Compared to the results reported for this database in its corresponding article, TattTRN improves identification performance by 18% (*i.e.*, 81.60% vs. 64%). Despite the results obtained by the TattTRN, it still fails to encode extremely low-quality images or images containing overlapping tattoos. In future work, we will combine TattTRN with a new prompt-based component that includes the semantic meaning of tattoos and evaluate the effect of different skin tones on tattoo segmentation and retrieval.

## 7. Acknowledgement

This research work has been partially funded by the Hessian Ministry of the Interior and Sport in the course of the Bio4ensics project and the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.



## References

- [1] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1578–1587, 2022. 4
- [2] J. Calmon, J. Queiroz, C. Goes, and A. Loula. Augmented tattoo: Evaluation of an augmented reality system for tattoo visualization. In *Proc. Conf. on Graphics, Patterns and Images*, pages 265–272, 2015. 3
- [3] R. da Silva and H. Lopes. A transfer learning approach for the tattoo detection problem. In *Congresso Brasileiro de Inteligência Computacional*, pages 1–4, 2021. 1
- [4] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009. 5
- [5] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4690–4699, 2019. 4
- [6] X. Di and V. M. Patel. Deep tattoo recognition. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 51–58, 2016. 2
- [7] X. Di and V. M. Patel. Deep learning for tattoo recognition. *Deep Learning for Biometrics*, pages 241–256, 2017. 2
- [8] L. J. Gonzalez-Soler, C. Rathgeb, and D. Fischer. Semi-synthetic data generation for tattoo segmentation. In *Proc. Intl. Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2023. 1, 2, 3, 4, 5
- [9] L. J. Gonzalez-Soler, C. Rathgeb, D. Fischer, and K. M. Zyla. On the impact of tattoos on hand recognition. In *Proc. Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5, 2023. 4, 5
- [10] H. Han and A. Jain. Tattoo based identification: Sketch to image matching. In *Proc. Intl. Conf. on Biometrics (ICB)*, pages 1–8, 2013. 1, 2
- [11] H. Han, J. Li, A. K. Jain, S. Shan, and X. Chen. Tattoo image search at scale: Joint detection and compact representation learning. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 41(10):2333–2348, 2019. 1, 2, 5, 6, 8
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 4, 6
- [13] A. Howard, M. Sandler, G. Chu, L. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, et al. Searching for mobilenetv3. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1314–1324, 2019. 6
- [14] T. Hrkač, K. Brkić, and Z. Kalafatic. Tattoo detection for soft biometric de-identification based on convolutional neural networks. In *Proc. OAGM-ARW Joint Workshop*, pages 131–138, 2016. 1, 2
- [15] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7132–7141, 2018. 6
- [16] G. Huang, Z. Liu, L. Van Der Maaten, and K. Weinberger. Densely connected convolutional networks. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017. 6
- [17] M. Ibsen, C. Rathgeb, P. Drozdowski, and C. Busch. Face beneath the ink: Synthetic data and tattoo removal with application to face recognition. *Applied Sciences*, 12(24):12969, 2022. 4
- [18] M. Ibsen, C. Rathgeb, P. Drozdowski, and C. Busch. Face beneath the ink: Synthetic data and tattoo removal with application to face recognition. *Applied Sciences*, 12(24), 2022. 3
- [19] ISO/IEC JTC1 SC37 Biometrics. *ISO/IEC 19795-1:2021. Information Technology – Biometric Performance Testing and Reporting – Part 1: Principles and Framework*. International Organization for Standardization, 2021. 2, 5
- [20] A. Jain, J. Lee, and R. Jin. Tattoo-id: Automatic tattoo image retrieval for suspect and victim identification. In *Proc. Advances in Multimedia Information Processing (PCM)*, pages 256–265, 2007. 2
- [21] I. Joshi, M. Grimmer, C. Rathgeb, C. Busch, F. Bremond, and A. Dantcheva. Synthetic data in human analysis: A survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2024. 2, 3
- [22] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of stylegan. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8110–8119, 2020. 4
- [23] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 4
- [24] J. Lee, A. Jain, W. Tong, et al. Image retrieval in forensics: tattoo image database application. *IEEE MultiMedia*, 19(1): 40–49, 2011. 1, 2
- [25] B. Li, K. Xue, B. Liu, and Y. Lai. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1952–1961, 2023. 4
- [26] F. Li, W. Tong, R. Jin, A. Jain, and J. Lee. An efficient key point quantization algorithm for large scale image retrieval. In *Proc. First ACM workshop on Large-scale multimedia Retrieval and Mining*, pages 89–96, 2009. 1, 2
- [27] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10012–10022, 2021. 6
- [28] A. Makrushin, C. Kauba, S. Kirchgasser, S. Seidlitz, C. Kraetzer, A. Uhl, and J. Dittmann. General requirements on synthetic fingerprint images for biometric authentication and forensic investigations. In *Proc. ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec)*, pages 93–104, 2021. 2
- [29] R. McCabe. Information technology: American national standard for information systems: Data format for the interchange of fingerprint, facial, & scar mark & tattoo (smt) information. 2000. 2

- [30] J. Mun, K. Janigo, and K. Johnson. Tattoo and the self. *Clothing and Textiles Research Journal*, 30(2):134–148, 2012. [1](#)
- [31] M. Ngan and P. Grother. Tattoo recognition technology-challenge (tatt-c): an open tattoo database for developing tattoo recognition research. In *Proc. Intl. Conf. on Identity, Security and Behavior Analysis (ISBA 2015)*, pages 1–6, 2015. [1](#)
- [32] M. Ngan, G. W. Quinn, and P. Grother. Tattoo recognition technology-challenge (tatt-c) outcomes and recommendations: Revision 1.0, 2016. [1](#)
- [33] M. Nicolás-Díaz, A. Morales-González, and H. Méndez-Vázquez. Deep generic features for tattoo identification. In *Proc. Iberoamerican Congress on Pattern Recognition (CIARP)*, pages 272–282, 2019. [1](#), [2](#), [3](#), [5](#), [6](#)
- [34] M. Nicolás-Díaz, A. Morales-González, and H. Méndez-Vázquez. Weighted average pooling of deep features for tattoo identification. *Multimedia Tools and Applications*, 81(18):25853–25875, 2022. [2](#), [3](#), [6](#)
- [35] A. Nurhudatiana and A. Kong. On criminal identification in color skin images using skin marks (rppvsm) and fusion with inferred vein patterns. *IEEE Trans. on Information Forensics and Security ((TIFS))*, 10(5):916–931, 2015. [5](#)
- [36] A. Nurhudatiana, A. Kong, L. Altieri, and N. Craft. Automated identification of relatively permanent pigmented or vascular skin marks (rppvsm). In *Proc. Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2984–2988, 2013. [5](#)
- [37] A. Nurhudatiana, A. Kong, K. Matinpour, D. Chon, L. Altieri, S. Cho, and N. Craft. The individuality of relatively permanent pigmented or vascular skin marks (rppvsm) in independently and uniformly distributed patterns. *IEEE Trans. on Information Forensics and Security ((TIFS))*, 8(6):998–1012, 2013. [5](#)
- [38] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019. [5](#)
- [39] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. [4](#)
- [40] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proc. Intl. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015. [4](#)
- [41] Z. Sun, J. Baumes, P. Tunison, M. Turek, and A. Hoogs. Tattoo detection and localization using region-based deep learning. In *Proc. Intl. Conf. on Pattern Recognition (ICPR)*, pages 3055–3060, 2016. [1](#)
- [42] M. Tan and Q. Le. Efficientnetv2: Smaller models and faster training. In *Proc. Intl. Conf. on Machine Learning (ICML)*, pages 10096–10106, 2021. [6](#), [7](#)
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. [6](#)
- [44] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5265–5274, 2018. [4](#)
- [45] Q. Xu, S. Ghosh, X. Xu, Y. Huang, and A. Kong. Tattoo detection based on cnn and remarks on the nist database. In *Proc. Intl. Conf. on Biometrics (ICB)*, pages 1–7, 2016. [1](#), [2](#)
- [46] S. Yun, D. Han, S. Oh, S. Chun, J. Choe, and Y. Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 6023–6032, 2019. [7](#)
- [47] L. Zhang, Z. He, Y. Yang, L. Wang, and X. Gao. Tasks integrated networks: Joint detection and retrieval for image search. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 44(1):456–473, 2020. [2](#)
- [48] J. Zhu, T. Park, P. Isola, and A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2223–2232, 2017. [4](#)