Chapter 18

Detecting Soft-Biometric Privacy Enhancement

Peter Rot, Peter Peer, and Vitomir Štruc

Abstract With the proliferation of facial analytics and automatic recognition technology that can automatically extract a broad range of attributes from facial images, so-called *soft-biometric privacy-enhancing techniques* have seen increased interest from the computer vision community recently. Such techniques aim to suppress information on certain soft-biometric attributes (e.g., age, gender, ethnicity) in facial images and make unsolicited processing of the facial data infeasible. However, because the level of privacy protection ensured by these methods depends to a significant extent on the fact that privacy-enhanced images are processed in the same way as non-tampered images (and not treated differently), it is critical to understand whether privacy-enhancing manipulations can be detected automatically. To explore this issue, we design a novel approach for the detection of privacy-enhanced images in this chapter and study its performance with facial images processed by three recent privacy models. The proposed detection approach is based on a dedicated attribute recovery procedure that first tries to restore suppressed soft-biometric information and based on the result of the restoration procedure then infers whether a given probe image is privacy enhanced or not. It exploits the fact that a selected attribute classifier generates different attribute predictions when applied to the privacy-enhanced and attribute-recovered facial images. This **prediction mismatch (PREM)** is, therefore, used as a measure of privacy enhancement. In extensive experiments with three popular face datasets we show that the proposed PREM model is able to accurately detect privacy enhancement in facial images despite the fact that the technique requires no supervision, i.e., no examples of privacy-enhanced images are needed for training.

P. Rot (✉) · V. Štruc

Faculty of Electrical Engineering, University of Ljubljana, Tržaška cesta 25,
1000 Ljubljana, Slovenia
e-mail: peter.rot@fe.uni-lj.si

V. Štruc

e-mail: vitomir.struc@fe.uni-lj.si

P. Peer

Faculty of Computer and Information Science, University of Ljubljana, Večna pot 113,
1000 Ljubljana, Slovenia
e-mail: peter.peer@fri.uni-lj.si

© The Author(s) 2022

C. Rathgeb et al. (eds.), *Handbook of Digital Face Manipulation and Detection*,
Advances in Computer Vision and Pattern Recognition,
https://doi.org/10.1007/978-3-030-87664-7_18

391

18.1 Introduction

Recent advances in computer vision, machine learning, and artificial intelligence have pushed the capabilities of automated recognition technology far beyond of what was possible only a few years ago [1–4]. Using state-of-the-art recognition techniques it is possible today to reliably link facial images to individuals and to infer a wide variety of (soft-biometric) attributes, such as gender, age, ethnicity, kin relations, or even health-related attributes, from images captured in less than ideal conditions [5–8]. These advances have made it possible to deploy face recognition technology across a number of application domains ranging from security, border control, and criminal investigations to entertainment, mobile gadgets, social media, autonomous driving, or even health services [9]. While the outlined developments have brought about many societal benefits, increased security, and made a multitude of everyday tasks considerably more convenient, the increased proliferation of face recognition techniques also resulted in privacy concerns related to the possible (mis)use of biometric (facial) data.

Driven by these concerns, a considerable amount of research is currently looking at privacy mechanisms that can provide a trade-off between the utility of the data for facial analytics on the one hand, and the privacy of individuals, on the other [10–13]. To comply with GDPR’s minimization principle,¹ facial analytics needs to limit the processing of information only to what is necessary in relation to the key utility of the system. For instance, in face verification systems, where the key utility is to validate an identity claim and user consent is typical only given for this specific use case, automatic inference of potentially sensitive soft-biometric information should not be possible. Nevertheless, recent research shows [14] that a multitude of sensitive information can still be extracted from the data processed within common face verification systems. This information can potentially be misused, without user’s consent for other purposes (i.e., function creep) such as automatic targeted advertising, user profiling, or discrimination.

Existing mechanism used for ensuring privacy with facial images are usually based on deidentification technology [15, 16]. Such technology aims to conceal (suppress, remove, or replace) potentially sensitive visual information in images with the goal of privacy protection and can broadly be categorized into two distinct groups: (i) techniques that target identity, and (ii) techniques that focus on soft-biometric information. Solutions from the first group are useful for privacy protection when sharing visual data captured by third parties on various services, e.g., Google StreetView, where people may appear in the captured data, or when analyzing surveillance footage to protect the privacy of innocent bystanders. Solutions from the latter group also referred to as *deidentification techniques for soft-biometric identifiers* [15, 17] or (more recently) *soft-biometric privacy-enhancing techniques* [10, 18] are relevant, e.g., in the context of social media, where people are in general willing to share their

¹ GDPR Data Minimization Principle: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/principles/data-minimisation/#data_minimisation.

images online with friends and families, but typically object to privacy intrusions (e.g., targeted ads) facilitated by automatic processing of the uploaded images. In such settings, soft-biometric privacy-enhancing techniques (or soft-biometric privacy models) that try to manipulate facial images in a way that makes automatic extraction of facial attributes, such as age, gender, or ethnicity, challenging but preserves the visual appearances of the input images as much as possible are highly desirable. Such techniques are also at the heart of this chapter.

While a number of soft-biometric privacy models have been proposed in the literature over the years, many still rely (to some extent) on the concept of *privacy through obscurity*, where (improved) privacy protection is ensured as long as the privacy-enhanced images are processed in the same way as all others. If a potential adversary launches a reconstruction attack and tries to recover the suppressed attribute information, the privacy protection may be rendered ineffective [19]. It is, therefore, critical to understand to what extent privacy enhancement can be detected. If an adversary is able to detect that an image has been tampered with, he/she may use specialized analysis tools, manual inspection, or other more targeted means of inferring the concealed information. The *detectability* of privacy enhancement is, hence, a key aspect of existing privacy models that to a great extent determines the level of privacy protection ensured by privacy-enhanced images in real-world settings. However, despite its importance and implications for the deployment of soft-biometric privacy models in real-world applications, this issue has not yet been explored in the open literature.

In this chapter we try to address this gap and present a study centered around the task of detecting image manipulations caused by soft-biometric privacy models. Our goal is (i) to assess whether privacy enhancement can be detected automatically, and as a result (ii) to provide insight into privacy risks originating from such detection techniques. To facilitate the study, we develop a novel detection approach that uses a super-resolution based procedure to first recover suppressed attribute information from privacy-enhanced facial images and then exploits the PREDiction Mismatch (PREM) of an attribute classifier applied to facial images before and after attribute recovery to flag privacy-enhanced data. The proposed approach is evaluated in extensive experiments involving three recent privacy models and three public face datasets. Experimental results show that PREM is not only able to detect privacy enhancement with high accuracy across different data characteristics and privacy models used, but also that it ensures highly competitive results compared to related detection techniques from the literature.

In summary, we make the following key contributions in this chapter:

- We introduce, to the best of our knowledge, the first technique for the detection of soft-biometric privacy enhancement in facial images. The proposed technique, called PREM, measures the Kullback–Leibler divergence between the predictions of a soft-biometric attribute classifier applied to facial images before and after attribute recovery. As we discuss in the chapter, PREM (i) exhibits several desirable characteristics, (ii) requires no examples of privacy-enhanced images for training, and (iii) is applicable under minimal assumptions.

- We show, for the first time, that it is possible to detect privacy-enhanced facial images with high accuracy across a number of datasets and privacy models, suggesting that the *detectability* of privacy-related image tampering techniques represents a major privacy risk.
- We demonstrate the benefit of designing PREM in a learning-free manner through comparative experiments with a related detection technique from the literature.

The rest of the chapter is structured as follows: Sect. 18.2 provides relevant background information, presents an overview of the state-of-the-art in soft-biometric privacy-enhancing techniques, and elaborates on the importance of detecting face manipulations caused by privacy enhancement. Section 18.3 describes PREM and discusses its characteristics. Section 18.4 presents the experimental setup used for the evaluation of the proposed detection approach and discusses results and findings. Section 18.5 concludes the chapter with a discussion of the key findings and directions for future work.

18.2 Background and Related Work

This section provides background information on the topic of soft-biometric privacy enhancement and reviews relevant prior work. A more comprehensive review of the broader field of visual privacy and advances in the area of privacy protection with facial images is presented in some of the excellent recent surveys on these topics, e.g., [10, 15, 17, 20, 21].

18.2.1 Problem Formulation and Existing Solutions

Soft-biometric privacy enhancement can formally be defined as follows: given an original face image, $I \in \mathbb{R}^{w \times h}$, where w and h are the image width and height in pixels, and a soft-biometric attribute classifier ξ_a , where

$$\xi_a : I \mapsto \{a_1, a_2, \dots, a_N\}, \quad (18.1)$$

and the attribute labels $\{a_i\}_{i=1}^N$ correspond to classes $\{C_1, C_2, \dots, C_N\}$, the goal of soft-biometric privacy enhancement, ψ , is to generate privacy-enhanced images,

$$I_{pr} = \psi(I) \in \mathbb{R}^{w \times h}, \quad (18.2)$$

from which ξ_a cannot correctly predict the class labels a_i . Because the goal of ψ is to conceal selected soft-biometric information from automatic classification techniques without significantly altering the visual appearance of the images for human observers, an additional constraint is commonly considered when designing ψ , that

is, that the privacy-enhanced images *have to be* as close as possible to the originals, i.e., $I_{pr} \approx I_{or}$, in terms of some target objective, for example, the Mean Squared Error (MSE) or Structural Similarity (SSIM). The goal of this constraint is to preserve the utility of the data after privacy enhancement, e.g., to ensure that the privacy-enhanced images can be shared with friends and family on the web, while making unsolicited automatic processing infeasible.

From a conceptual point of view, existing soft-biometric privacy-enhancing techniques can be categorized into the following two groups depending on whether they try to induce:

- **Misclassifications:** Solutions from the first group typically rely on adversarial perturbations (and related strategies) and enhance privacy by inducing misclassifications, i.e., $\xi_a(I_{pr}) \neq \xi_a(I_{or})$, where an incorrect attribute class label is predicted from I_{pr} with high probability.
- **Equal class probabilities:** Solutions from the second group most often rely on (input-conditioned) synthesis models that enhance privacy by altering image characteristics in such a way that equal class posteriors are generated by the considered attribute classifier ξ_a given a privacy-enhanced image I_{pr} , i.e., $p(C_1|I_{pr}) \approx p(C_2|I_{pr}) \approx \dots \approx p(C_N|I_{pr})$.

18.2.2 Soft-Biometric Privacy Models

A considerable amount of research has been conducted on the topic of soft-biometric privacy over recent years [12, 13, 18, 22–25]. Mirjalili and Ross [26], for instance, presented a privacy-enhancing technique that suppresses gender information in face images. The technique applies Delaunay triangulation over prominent facial landmarks to represent input faces as a set of triangles. In the next step, the texture within these triangles is modified in such a way that a targeted gender classifier produces unreliable classification results, while the original texture appearance is preserved as much as possible. The authors showed that such an approach leads to image perturbations that efficiently suppress gender information but have only a minimal impact on visual appearance and in turn on verification accuracy.

Another approach to soft-biometric gender privacy was presented in [27]. Here, the authors proposed a so-called Semi-Adversarial Network (SAN) that utilizes *conditional image synthesis* to suppress gender information in facial images. SAN models represent convolutional auto-encoders that are trained in an adversarial setting using two discriminators, where the first aims to enforce gender privacy and the second tries to retain verification accuracy (i.e., image similarities). The results reported by the authors show that the SAN model is able to obscure a high degree of gender information in facial images, while retaining the utility of the data for biometric verification purposes. An extension of this work was also presented later with the goal of improving the generalization capabilities of the initial SAN to unseen classifiers [23]. The main idea of the extended FlowSAN model is to utilize multiple

SAN transforms successively with the aim of making the privacy enhancement less dependent on a single target (gender) classifier. FlowSAN models were shown to offer better generalization capabilities than the simpler one-stage SANs, while still offering a competitive trade-off between privacy protection and utility preservation.

To make SAN models applicable beyond the (binary) problem of gender privacy, Marialli et al. [12] proposed PrivacyNet [12], an advanced SAN model based on the concept of Generative Adversarial Networks (GANs). Different from previous techniques aimed at soft-biometric privacy, PrivacyNet was demonstrated to be capable of privacy enhancement with respect to multiple facial attributes, including race, age and gender. PrivacyNet, hence, generalized the concept of SAN-based privacy enhancement to arbitrary combinations of soft-biometric attributes. The SAN-family of algorithm falls into the second group of techniques discussed above and tries to induce equal class probabilities with the considered attribute classifiers.

A misclassification-based approach to privacy enhancement of k facial attributes via adversarial perturbations (k -AAP) was introduced by Chhabra et al. in [11]. The approach aims to infuse facial image with so-called adversarial noise with the goal of suppressing a predefined set of arbitrary selected facial attributes, while preserving others. k -AAP is based on the Carlini Wagner L2 attack [28] and achieves promising results with attribute classifiers that were included in the design of the adversarial noise. The approach results in image perturbations (for most images) that are efficient at obscuring attribute information for machine learning models but are virtually invisible to a human observer. However, similarly as the original version of SAN, k -AAP does not generalize well to arbitrary classifiers. The idea of flipping facial attributes using adversarial noise was also explored in [22] where the authors investigated the robustness of facial features perturbations generated by the Fast Gradient Sign Method (FGSM) method [29].

While the techniques reviewed above try to manipulate facial images to ensure soft-biometric privacy, recent research is also looking at techniques that try to modify image representation and suppress attribute information at the representation (or feature) level [13, 18, 30–32]. However, such techniques are based on different assumptions and relevant mostly in the context of biometric systems. In this work we, therefore, focus only on the more general topic of image-level soft-biometric privacy enhancement.

18.2.3 Detecting Privacy Enhancement

When evaluating soft-biometric privacy-enhancing techniques, the literature typically focuses on performance and the level of privacy protection the techniques can offer. Other aspects are commonly of less interest and, as a results, are significantly less explored. This leads to several interesting research questions, e.g., *To what extent can privacy enhancement be detected? Is it possible to flag facial images that have been tampered with by privacy-enhancing techniques?* While a considerable body of work has been presented in the literature to detect traditional image tampering,

e.g., [33–35], existing detection methods have mostly been investigated within the digital forensics community. The problem of detecting privacy enhancement, on the other hand, has not yet been explored in the literature. Because some privacy models are based on adversarial perturbations, this problem is also partially related to adversarial-attack detection techniques [36–40]. However, since soft-biometric privacy enhancement also includes synthesis-based methods, data hiding solutions and a wide variety of other approaches [10], the problem of detecting such image modifications is considerably broader. In the remainder of the chapter, we present our solution to the problem of privacy enhancement detection, which relies on a simple attribute recovery procedure.

18.3 Tampering Detection Through Prediction Mismatch (PREM)

In this section we now describe the proposed approach for detecting privacy enhancement (or tampering) in facial images. We focus on gender experiments in this chapter, but the proposed approach is general and can be applied to arbitrary soft-biometric attributes.

18.3.1 PREM Overview

Soft-biometric privacy enhancement aims to introduce minute changes into facial images in such a way that the predictions of a selected attribute classifier become unreliable, while the original appearance of the facial images is preserved as much as possible. As a result, existing soft-biometric privacy models commonly add (targeted) high-frequency components to the input images, which are imperceptible to human observers, but adversely affect the performance of automatic recognition techniques. Based on these characteristics, we design a detection technique in this chapter that tries to identify images that were tampered with by soft-biometric privacy models.

The main idea of our detection technique is illustrated in Fig. 18.1. At the core of the technique is an attribute recovery procedure, $\chi(\cdot)$, that tries to restore suppressed attribute information from the given input image I . For unaltered facial images a selected soft-biometric attribute classifier ζ is expected to produce similar predictions before and after attribute recovery. For tampered images, on the other hand, a mismatch is expected between the predictions generated from the input image I and the corresponding attribute-recovered version I_{re} . This **PRE**diction **M**ismatch (**PREM**) can then be exploited for detecting privacy enhancement. Details on PREM are provided in the following section.

Fig. 18.1 Overview of the proposed PREM techniques for the detection of soft-biometric privacy enhancements in facial images. At the core of the detection technique is a super-resolution (SR) based procedure that aims to restore suppressed attribute information. PREM exploits the mismatch between the predictions of a selected attribute (*gender* in this work) classifier generated from the original and attribute-recovered images to detect privacy enhancement. PREM is *learning-free* and, hence, requires no examples of privacy-enhanced images

18.3.2 Super-Resolution for Attribute Recovery

To restore soft-biometric attributes from privacy-enhanced facial images, we consider a simple super-resolution (or hallucination²) approach. Super-resolution can be seen as a special type of restoration approach that selectively adds specific details (high-frequency components) to low-resolution input images. Because the high-frequency details are added in a selective manner based on learnt correspondences between pairs of low- and high-resolution facial images, images subjected to super-resolution are in essence remapped to a higher resolution, which impacts image characteristics, including those infused by privacy enhancement. The use of super-resolution for attribute recovery is further motivated by its success in mitigating the impact of adversarial noise, which is also used as a privacy-enhancing mechanism by some of the existing privacy models.

We propose a straight-forward two-step procedure to remove the effect of privacy enhancement on the appearance and characteristics of facial images—see Fig. 18.2 for an illustration. In the first step, the input image, $I_{pr} \in \mathbb{R}^{w \times h}$ is downsampled by a factor s , i.e.,

$$I_{pr}^{(s)} = f_s(I_{pr}), \quad (18.3)$$

where $f_s(\cdot)$ is a simple bilinear down-sampling function conditioned on s and $I_{pr}^{(s)} \in \mathbb{R}^{w_s \times h_s}$ is the downsampled image with $w_s < w$ and $h_s < h$. This down-sampling step acts as a low-pass filter that removes high-frequency information from the image and is, therefore, expected to impact (high-frequency) image artifacts introduced by the privacy models the most. In the second step, the procedure then restores the

² Note that single image super-resolution techniques are referred to as face hallucination methods when applied to facial images. This term illustrates the fact the high-frequency components are hallucinated by the super-resolution model during up-sampling [41–43].

Fig. 18.2 Overview of the proposed approach to attribute recovery through super-resolution (SR). The privacy-enhanced image, I_{pr} is first downsampled to remove high-frequency components introduced by the privacy models. High-frequency components are then selectively added back into the restored image, I_{re} , through a super-resolution model

downsampled image to its original size ($w \times h$) and recovers semantically meaningful image details (with attribute information) through super-resolution-based upscaling, i.e.,

$$I_{re} = h_s(I_{pr}^{(s)}) \in \mathbb{R}^{w \times h}, \quad (18.4)$$

where h_s is a super-resolution model that upscales images by a factor of s . The high-frequency information that is lost with the down-sampling procedure is, to a certain degree, reconstructed with the super-resolution model, such that the output image is similar to the input. However, as illustrated in Fig. 18.2 the super-resolved image is typically slightly smoother and contains hallucinated image details.

We denote the presented attribute recovery procedure as $\chi : I_{pr} \mapsto I_{re}$ and use the state-of-the-art C-SRIP model from [43] for our implementation. Here, I_{re} stands for the attribute-recovered images, shown on the right side of Fig. 18.2. We note at this point, that the described recovery procedure is model agnostic, so any super-resolution model could be used instead of C-SRIP with similar results. C-SRIP was selected for our implementation because of its state-of-the-art performance and the fact that it is publicly available.

18.3.3 Measuring the Prediction Mismatch

As suggested above, the proposed PREM detection technique exploits the fact that for privacy-enhanced images, I_{pr} , a selected attribute classifier ξ generates different posterior probabilities $p(C_k|I_{pr})$ than for images subjected to the presented super-resolution based attribute recovery procedure, $p(C_k|I_{re})$. Because we consider the task of detecting privacy enhancements aimed at suppressing gender information in facial images, similarly to [18, 23, 44, 45], the classes C_k are defined as $C_k \in \{C_m, C_f\}$, where m denotes the male and f the female class. However, we note that in general the same conceptual solution could be applied to arbitrary attribute classes C_k . By comparing the posteriors, $p(C_k|I_{pr})$ and $p(C_k|I_{re})$, it is possible

Algorithm 1: Detecting soft-biometric privacy enhancement with PREM

Input: Input probe image I
Output: Detection score τ

- 1 Implement attribute recovery procedure χ based on SR model h_s —Eqs. (18.3) and (18.4);
 - 2 Attempt attribute recovery: $I_{re} = \chi(I)$;
 - 3 Use ξ for classification over I and I_{re} ;
 - 4 Calculate detection score $\tau = D_{SKL}$ —Eq. (18.5);
 - 5 Determine whether I is tampered based on τ ;
-

to determine whether an image has been tampered with or not. In this work we utilize a symmetric version of the Kullback–Leibler divergence [46] to compare the distributions generated by the attribute classifier, i.e.,

$$D_{SKL}(p, q) = D_{KL}(p||q) + D_{KL}(q||p), \quad (18.5)$$

where

$$D_{KL}(p||q) = \sum_{x \in \mathcal{X}} (p(x) - q(x)) \log \left(\frac{p(x)}{q(x)} \right), \quad (18.6)$$

where $p = p(C_k|I_{pr})$, $q = p(C_k|I_{re})$, $\mathcal{X} = C_k$. To facilitate the detection of soft-biometric privacy enhancement (or tampering), D_{SKL} is used as a detection score τ in our experiments. A decision whether the probe image I was tampered with or not can be made based on the value of τ . D_{SKL} is in general bounded by $[0, \infty]$, where a value of 0 indicates that there is no difference between distributions.

18.3.4 PREM Summary and Characteristics

A high-level summary of the proposed PREM detection technique is given in Algorithm 1. The techniques have several desirable characteristics, i.e.,

- **Training free detection:** Unlike competing techniques for image tampering detection, PREM is training-free and does not require any examples of privacy-enhanced images for training. PREM relies solely on the fact that the output of a (in this case gender) classifier changes after attribute recovery compared to the output produced with the original face image. Thus, the detection scheme is expected to work, across a wide range of privacy enhancement techniques as long as they are based on the same assumptions, i.e., minimal changes in appearance compared to the original images.
- **Complementarity to supervised detection techniques:** While the task of detecting privacy enhancement in facial images is new, there are related tampering-detection techniques that could be adopted for this task. Because these techniques

typically rely on supervision [40, 47], PREM provides complementary information that can be combined with supervised techniques to further improve performance.

- **Generality:** The concept exploited by PREM is not limited to the described super-resolution based recovery approach and can be implemented with any techniques that is able to restore suppressed soft-biometric attributes. Thus, PREM can easily accommodate more advance recovery schemes and is expected to benefit from future advances in this problem area.

18.4 Experiments and Results

The experiments, presented in this section, aim at (i) evaluating the extent to which soft-biometric privacy enhancement can be detected, (ii) analyzing the performance of PREM in blue comparison to alternatives from the literature, (iii) exploring the complementarity of PREM and competing detection techniques, and (iv) providing insight into the behavior of the technique and its limitations. As emphasized earlier, we focus on privacy-enhancing techniques that are trained to suppress *gender information*, which is also the most frequent attribute considered in the literature when studying soft-biometric privacy models [18, 23, 44, 45].

18.4.1 Datasets and Experimental Setup

We use three publicly available face datasets for the experiments, i.e., Labeled Faces in the Wild (LFW) [48], MUCT [49], and Adience [50]. The datasets come with the necessary attribute labels (for gender in this case) and contain facial images captured in unconstrained settings and in a wide variety of imaging conditions—a few examples are shown in Fig. 18.3. Additionally, the datasets are commonly used in research on soft-biometric privacy enhancement, e.g., [12, 18, 23, 30]) and are, therefore, also selected for this work.

To have a consistent starting point for the experiments, all images are roughly aligned, so that faces are approximately centered. Next, the facial region is cropped to exclude background pixels and finally, the images are rescaled to a standard size of 224×224 pixels. Because gender classifiers are needed to train the soft-biometric privacy models, one such classifier ζ is trained for each dataset using a gender-balanced set of training images, i.e., around 5700 for LFW, 10400 for Adience and 3600 for MUCT. A separate set of images from each dataset is reserved for the testing stage. This testing set does not overlap in terms of subjects with the training data.

Fig. 18.3 Example images from the three datasets used in the experimental evaluation. The datasets were captured in challenging and diverse (real world) imaging conditions posing a significant challenge to both privacy enhancement as well as detection techniques

18.4.2 Utilized Privacy Models

We implement three recent soft-biometric privacy models for the experiments, i.e., the k -AAP approach from [11], the FGSM-based method from [51] and the FlowSAN technique from [23]. FlowSAN can in general be implemented with different trade-offs between the level of attribute suppression ensured and the preserved utility of the facial images. We, therefore, consider two different model variants in our experiments: (i) one with a sequence of three SAN models (i.e., FlowSAN-3), and (ii) one with a sequence of five SAN models (i.e., FlowSAN-5). The selected techniques are utilized for the evaluation because of their state-of-the-art performance and the fact that they rely on different privacy mechanisms. A few examples of the impact of the privacy models on the visual appearance of selected facial images are presented in Fig. 18.4. Note that the models result in visually different image manipulations. While the appearance of the images processed with the adversarial techniques k -AAP and FGSM is very close to the originals, the FlowSAN models introduce larger changes. However, these changes also help with the robustness of the models with respect to unseen gender classifiers. We note that the appearance of the FlowSAN enhanced images in Fig. 18.4 is a direct consequence of the design of the FlowSAN models, which can be implemented to provide a trade-off between visual quality and privacy protection—larger degradations ensure better privacy protection, while higher quality images result in lower levels of privacy.

18.4.3 Implementation Details

Gender classifiers. A VGG16 model architecture [52] is used to implement the gender classifiers ζ needed for the privacy enhancement. We adapt a pretrained VGGFace2 [53] for gender classification by replacing the original softmax layer

Fig. 18.4 Impact of the privacy models on the visual appearance of facial images. The adversarial methods (*k*-AAP and FGSM) introduce only minor changes, whereas the FlowSAN models generated more pronounced appearance differences—note that these models are designed to be robust to unseen classifiers, which is not the case with the adversarial techniques

Table 18.1 Values of hyper-parameters utilized for the implementation of the selected soft-biometric privacy models

Model	Strategy	Parameters
<i>k</i> -AAP	Fixed parameters	#Max iter.: 20, Learn. rate: 0.01, Init. const.: 1, 000
FGSM	Value search	Range: [0.9–0.51], #Epochs: 200, Epsilon: 0.5
FlowSAN	Two versions	#SANs in sequence: 3 and 5

with a two-class softmax. The models are then learned by fine tuning the last two layers using the Adam optimizer with a learning rate of 0.001, momenta $\beta_1 = 0.9$ and $\beta_2 = 0.999$ and an ϵ value of 10^{-7} . For the FlowSAN models several such classifiers are trained using different configurations of training data, as suggested in [54].

Super-resolution model. To implement the attribute recovery procedure, the C-SRIP³ [43] super-resolution model is selected. The model is trained on the Casia WebFace dataset [55]. For the recovery procedure, χ , images are first resized from the original size of 224×224 to 64×64 pixels and then subjected to the super-resolution model to upsample the images back to the initial image resolution.

Privacy models. The three soft-biometric privacy models selected for our experiments, *k*-AAP, FGSM and FlowSAN (FlowSAN-5) rely on several hyper-parameters. These are selected based on preliminary experiments and the suggestions made by the authors of the models. Table 18.1 provides a summary on the hyper-parameters used in the experiments.

³ C-SRIP: <https://lmi.fe.uni-lj.si/en/research/fh/>.

18.4.4 Results and Discussions

PREM Evaluation. In the first series of experiments, we analyze the performance of PREM with respect to the detection of privacy-enhanced images. To this end, we use a gender-stratified test set of 698 privacy-enhanced and 698 original images for each privacy model (k -APP, FGSM, FlowSAN-3 and FlowSAN-5) and each dataset considered in the experiments. We partition the test images into 4 disjoint data splits, over which we report results. Because privacy enhancement detection is framed as a two-class classification problem in this work, we report results graphically in terms of Received Operating Characteristics (ROC) curves and quantitatively in the form of the Area Under the ROC Curve (AUC). The 4 data splits are also utilized to generate confidence intervals for the ROC curves.

Figure 18.5 shows the average ROC curves and corresponding confidence intervals for this first experimental series. Because the performance differences across different privacy models are small, the curves are presented on a semi-log scale for better visualizations. AUC scores of the experiments are reported in Table 18.2. As can be seen, PREM achieves almost ideal detection performance for all privacy models on the LFW and MUCT datasets and with average AUC scores of 0.629 for k -AAP, 0.858 for FGSM, 0.775 for FlowSAN-3 and 0.793 for FlowSAN-5 also performs well on Adience. The slightly weaker results on Adience are a consequence of the challenging imaging characteristics present in the dataset and are in line with observations made in the open literature [18], where different models (addressing various tasks) perform worse on this dataset.

Comparison with Competing Models. Note again that the task of detecting soft-biometric privacy enhancement in facial images has not been explored widely in the

Fig. 18.5 ROC curves generated for the privacy enhancement detection experiments with the proposed PREM detection approach: **a** on LFW, **b** on MUCT and **c** on Adience. Note that the detection of the privacy enhancement is more successful on the LFW and MUCT, than on Adience, which is a highly challenging face dataset captured in real-world settings

Table 18.2 AUC scores ($\mu \pm \sigma$) generated in the detection experiments. The learning-free (black-box) PREM technique is compared against the recent (supervised) T-SVM approach to adversarial-attack detection from [40]. Results are reported for white-box as well as black-box settings

Priv. Model	Dataset	PREM (ours)	White-box	Black-box		Combined (ours)
			T-SVM	T-SVM (A) [†]	T-SVM (B) [‡]	
<i>k</i> -AAP	LFW	0.957 ± 0.005	0.984 ± 0.007	0.743 ± 0.013	0.984 ± 0.007	0.946 ± 0.020
	MUCT	0.965 ± 0.005	0.534 ± 0.021	0.727 ± 0.030	0.534 ± 0.021	0.939 ± 0.024
	Adience	0.629 ± 0.007	0.541 ± 0.007	0.604 ± 0.008	0.541 ± 0.007	0.638 ± 0.008
FGSM	LFW	0.989 ± 0.003	0.955 ± 0.011	0.894 ± 0.013	0.921 ± 0.016	0.988 ± 0.006
	MUCT	0.989 ± 0.003	0.852 ± 0.012	0.858 ± 0.015	0.552 ± 0.030	0.985 ± 0.007
	Adience	0.858 ± 0.003	0.624 ± 0.009	0.595 ± 0.016	0.471 ± 0.014	0.831 ± 0.013
FlowSAN-3	LFW	0.987 ± 0.002	1.000 ± 0.000	1.000 ± 0.000	0.997 ± 0.002	0.996 ± 0.004
	MUCT	0.993 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	0.524 ± 0.020	1.000 ± 0.000
	Adience	0.775 ± 0.002	1.000 ± 0.000	1.000 ± 0.000	0.609 ± 0.005	0.986 ± 0.003
FlowSAN-5	LFW	0.988 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	0.998 ± 0.001	0.996 ± 0.004
	MUCT	0.994 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	0.531 ± 0.018	1.000 ± 0.000
	Adience	0.793 ± 0.011	1.000 ± 0.000	1.000 ± 0.000	0.563 ± 0.005	0.988 ± 0.003

[†] T-SVM trained on FlowSAN-5 and training data from LFW[‡] T-SVM trained on *k*-AAP and training data from LFW

literature. However, because privacy enhancement techniques share some characteristics with adversarial attacks, we select the recent state-of-the-art transformation-based detection technique T-SVM from [40] as a baseline for the experiments and compare it to the proposed PREM approach. T-SVM utilizes a combination of the discrete wavelet transform (DWT) and the discrete sine transform (DST) to represent input images and relies on a support vector machine (SVM) for privacy enhancement detection. As such, it requires training data to be able to distinguish between original and privacy-enhanced images. We, therefore, consider two distinct experimental settings in the evaluations to compare PREM and T-SVM, i.e.,

- **White-box setting:** In this setting, unfiltered access to all privacy models is assumed. Thus, T-SVM is trained separately for each privacy model using corresponding training images. The training part of LFW (i.e., around 5700 images) is utilized to learn the detection model and the test part of LFW, MUCT and Adience is used for the evaluation.
- **Black-box setting:** In this setting, access to some privacy-enhanced images is assumed, but not to images of all considered privacy models. As a result, T-SVM needs to generalize to unseen privacy models when trying to detect tampering. Two different detection models are considered in the experiments. The first, T-SVM (A), is trained for the detection of FlowSAN-5 enhanced images and only on data from LFW. The second, T-SVM (B), is trained for detecting *k*-AAP based enhancement, again only on LFW. Both models are tested on all datasets and with all privacy models.

The results in Table 18.2 shows that in the white-box scenario T-SVM is comparable to PREM on LFW in terms of AUC scores for k -AAP and FGSM, but significantly worse on the MUCT and Adience datasets. This observation suggests that despite access to privacy-enhanced training images for all privacy models, the performance of T-SVM is affected by the data characteristics, which is much less the case with the proposed *learning-free* PREM technique. For the experiments with the FlowSAN models, T-SVM offers ideal detection performance on all three datasets. PREM, on the other hand, is very competitive on the images from LFW and MUCT, but performs weaker on images from the Adience dataset. The presented results suggest that the generalization ability of the white-box T-SVM technique is below PREM despite access to privacy-enhanced training examples for each privacy model, which are not required for PREM.

When examining the AUC scores for the black-box setting, we see that T-SVM performs somewhat weaker than in the white-box scenario. T-SVM (A), trained with FlowSAN-5 processed images, is still able to ensure ideal detection performance for both FlowSAN variants on all three experimental datasets, but is less competitive with images enhanced with k -AAP and FGSM. For these privacy models the detection performance of T-SVM (A) drops slightly behind the white-box T-SVM and considerably behind PREM. T-SVM (B), trained with k -AAP images of LFW, is the least competitive and only able to match the performance of PREM for a few selected cases involving LFW. These results clearly demonstrate the added value of training-free privacy enhancement detection, where the detection performance is not affected by the available training data, which otherwise considerably affects the generalization capabilities of the detection techniques.

Combined Privacy Enhancement Detection. In the previous section we demonstrated that PREM is able to detect privacy enhancement in facial images with high accuracy. We also observed that the competing method, T-SVM, is highly competitive on images where PREM performed weaker, e.g., with the FlowSAN models on Adience. In this section, we, therefore, try to make use of the best of both worlds and explore whether a combination of the two detection approaches can further improve results. To make the experimental setup as challenging as possible, we assume a black-box experimental setting for T-SVM (marked T-SVM (A) in Table 18.2) and train it with LFW images processed by FlowSAN-5. We then consider this model with all other (unseen) privacy models and all three datasets.

Given an input face image I , both PREM and T-SVM produce a detection score, $\tau_{prem}(I)$ and $\tau_{tsvm}(I)$, that is utilized to make a decision on whether images were tampered with or not. Since the two detection models are independent, we combine the computed detections scores using a product fusion rule of the following form:

$$\tau = \tau_{prem}^{\beta} \cdot \tau_{tsvm}^{1-\beta}, \quad (18.7)$$

where τ stands for the combined detection score and β represents a parameter that balances the relative contribution of the two scores and is set to $\beta = 0.5$ in this work.

The detection results for the combined approach are presented on the right part of Table 18.2. As can be seen, the combination of the detection scores leads to a

comparable detection performance to PREM for k -AAP and FGSM on images from the LFW and MUCT datasets. Here, PREM convincingly outperformed T-SVM in the comparative evaluation and the combined approach is able to match PREM's AUC score with these privacy models. For privacy enhancement with the FlowSAN models, the combination of the detection techniques again performs comparable to the better performing of the two individual models on all experimental datasets. Here, the combined detection approach achieves close to perfect performance (also on Adience) and further improves on the initial capabilities of PREM, which was the weaker of the two detection models on this dataset.

Overall, the results suggest that the combination of PREM and T-SVM is beneficial for the detection of privacy enhancement and contributes toward more robust performance across different privacy models and data characteristics. The conceptual differences between PREM and T-SVM allow the combined approach to capture complementary information for the detection procedure and produce more reliable detection scores across all experiments.

Visual analysis. In the last part of our evaluation, we conduct a qualitative analysis and investigate what data characteristics cause PREM to fail. A few examples of face images, where PREM generated errors at a decision threshold that ensures equal error rates (EERs) with the ROC curves from Fig. 18.5 is presented in Fig. 18.6. Here, images in the top row correspond to cases where PREM failed to detect privacy enhancement. As can be seen, these images contain visible image artifacts that cannot be recovered using the super-resolution in based attribute recovery procedure, resulting in minimal changes in the gender predictions and subsequently misdetections. Images in the bottom row represent examples that have been classified as being privacy enhanced but in fact represent original unaltered faces. In most cases,

Fig. 18.6 Visual examples of misdetections produced by PREM at a decision threshold corresponding to the equal error rate (ERR) on the ROC curves from Fig. 18.5. Images in the top row correspond to cases where PREM failed to detect the privacy enhancement and images in the bottom row correspond to examples incorrectly flagged by PREM as being privacy enhanced. Example results are reported for all privacy models considered in the experiments

these images are of poorer quality (due to blur, noise, etc.) and get improved by our attribute recovery procedure. As a result, the output of the gender classifier changes sufficiently to incorrectly classify the images as being privacy enhanced.

18.5 Conclusion

In this chapter we studied the problem of soft-biometric privacy enhancement and explored to what extent such privacy enhancement can be detected. To facilitate the study we designed a novel approach for the detecting privacy enhancements in facial images, called PREM. We evaluated PREM with three recent privacy enhancement techniques and three experimental datasets. Our experimental results showed that PREM was able to detect privacy enhancement with across all considered privacy models and facial images of very different characteristics. These findings have considerable implications for future research in privacy enhancement, e.g.,

- If privacy enhancement can be detected, the flagged images may be processed with alternative means and more elaborate analysis techniques. This type of processing could invalidate the effect of privacy enhancement.
- Privacy-enhancing techniques need to come with privacy guarantees to ensure that even after detection, privacy-enhanced images may not be misused. This points to the need for formal soft-biometric privacy models [10] that allow to quantify the level of privacy ensured. Such models have been considered with deidentification techniques [56] but have not yet been explored in the literature for soft-biometric privacy solutions.

As part of our future work, we plan to explore possibilities toward formal soft-biometric privacy models that address some of the challenges outlined above and offer privacy protection even if subjected to advanced means of processing.

Acknowledgements This research was supported in parts by the ARRS Research Project J2-1734 “Face Deidentification with Generative Deep Models”, ARRS Research Programmes P2-0250 (B) “Metrology and Biometric Systems” and P2-0214 (A) “Computer Vision”.

References

1. Guo G, Zhang N (2019) A survey on deep learning based face recognition. *Comput Vis Image Understand* 189:1–37
2. Ross A, Banerjee S, Chen C, Chowdhury A, Mirjalili V, Sharma R, Swearingen T, Yadav S (2019) Some research problems in biometrics: the future beckons. In: *International conference on biometrics (ICB)*, pp 1–8
3. Križaj J, Peer P, Štruc V, Dobrišek S (2019) Simultaneous multi-descent regression and feature learning for facial landmarking in depth images. *Neural Comput Appl*, pp 1–18
4. Batagelj B, Peer P, Štruc V, Dobrišek S (2021) How to correctly detect face-masks for COVID-19 from visual information? *Appl Sci* 11(5):1–24

5. Berthouze N, Valstar M, Williams, A, Egede J, Olugbade T, Wang C, Meng H, Aung M, Lane N, Song S (2020) Emopain challenge 2020: multimodal pain evaluation from facial and bodily expressions. In: IEEE conference on automatic face and gesture recognition (FG)
6. Puc A, Štruc V, Grm. K (2021) Analysis of race and gender bias in deep age estimation models. In: 28th European signal processing conference (EUSIPCO), pp 830–834
7. Gonzalez-Sosa E, Fierrez J, Vera-Rodriguez R, Alonso-Fernandez F (2018) Facial soft biometrics for recognition in the wild: Recent works, annotation, and cots evaluation. *IEEE Trans Inform Forensics Secur (TIFS)*, 13(8):2001–2014
8. Robinson JP, Shao M, Wu Y, Liu H, Gillis T, Fu Y (2018) Visual kinship recognition of families in the wild. *IEEE Trans Pattern Anal Mach Intell (TPAMI)* 40(11):2624–2637
9. Rattani A, Derakhshani R (2018) A survey of mobile face biometrics. *Comput Electr Eng* 72:39–52
10. Meden B, Rot P, Terhörst P, Damer N, Kuijper A, Scheirer W, Ross, A Peer P, Štruc. V (2021) Privacy-enhancing face biometrics: a comprehensive survey. Under Rev
11. Chhabra S, Singh, R Vatsa M, Gupta G (2018) Anonymizing k-facial attributes via adversarial perturbations. *International joint conferences on artificial intelligence (IJCAI)*
12. Mirjalili V, Raschka S, Ross A (2020) PrivacyNet: semi-adversarial networks for multi-attribute face privacy. *IEEE Trans Image Process (TIP)*, 29:9400–9412
13. Morales A, Fierrez J, Vera-Rodriguez R, Tolosana. R (2020) SensitiveNets: learning agnostic representations with application to face images. *IEEE Trans Pattern Anal Mach Intell (TPAMI)*
14. Philipp PT, Fährmann D, Damer N, Kirchbuchner F, Kuijper A (2020) Beyond identity: what information is stored in biometric face templates? In: 2020 IEEE international joint conference on biometrics (IJCB), pp 1–10
15. Ribarić S, Ariyaeeinia A, Pavesić N (2016) De-identification for privacy protection in multimedia content: a survey. *Signal Process: Image Commun* 47:131–151
16. Meden B, Malli RC, Fabijan S, Ekenel HK, Štruc V, Peer P (2017) Face Deidentification with generative deep neural networks. *IET Signal Process* 11(9):1046–1054
17. Garfinkel SL (2015) De-identification of Personal Information. National Institute of standards and technology (NIST)
18. Bortolato B, Ivanovska M, Rot P, Križaj J, Terhörst P, Damer N, Peer P, Štruc V (2020) Learning privacy-enhancing face representations through feature disentanglement. In: IEEE international conference on automatic face and gesture recognition (FG)
19. Rot P, Peer P, Struc V (2021) PrivacyProber: assessment and detection of soft-biometric privacy-enhancing techniques. Under Rev, 1–18
20. Winkler T, Rinner B (2014). Security and privacy protection in visual sensor networks: a survey. *ACM Comput Surv* 47(1)
21. Padilla-López JR, Chaaraoui AA, Flórez-Revuelta F (2015). visual privacy protection methods: a survey. *Expert Syst Appl* 42(9):4177–4195
22. Rozsa A, Günther M, Rudd EM, Boulton T (2019) Facial attributes: accuracy and adversarial robustness. *Pattern Recogn Lett* 124:100–108
23. Mirjalili V, Raschka S, Ross A (2019). FlowSAN: privacy-enhancing semi-adversarial networks to confound arbitrary face-based gender classifiers. *IEEE Access* 7:99735–99745
24. Terhörst P, Damer N, Kirchbuchner F, Kuijper A (2019) Unsupervised privacy-enhancement of face representations using similarity-sensitive noise transformations. *Appl Intell*, 1–18
25. Terhörst P, Huber M, Damer N, Rot P, Kirchbuchner F, Struc V, Kuijper A (2020) Privacy evaluation protocols for the evaluation of soft-biometric privacy-enhancing technologies. In: International conference of the biometrics special interest group (BIOSIG)
26. Mirjalili V, Ross A (2017) Soft biometric privacy: retaining biometric utility of face images while perturbing gender. In: International joint conference on biometrics (IJCB), pp 564–573
27. Mirjalili V, Raschka S, Namboodiri A, Ross A (2018) Semi-adversarial networks: convolutional autoencoders for imparting privacy to face images. In: International conference on biometrics (ICB), pp 82–89
28. Carlini N, Wagner D (2017) towards evaluating the robustness of neural networks. In: IEEE Symposium on security and privacy (SP), pp 39–57

29. Goodfellow I, Shlens J, Szegedy C (2015) Explaining and harnessing adversarial examples. In: International conference on learning representations
30. Terhöst P, Riehl K, Damer N, Rot P, Bortolato B, Kirchbuchner F, Štruc V, Kuijper A (2020) PE-MIU: a training-free privacy-enhancing face recognition approach based on minimum information units. *IEEE Access*
31. Terhöst P, Huber M, Damer N, Kirchbuchner F, Kuijper A (2020) Unsupervised enhancement of soft-biometric privacy with negative face recognition. *arXiv preprint [arXiv:2002.09181](https://arxiv.org/abs/2002.09181)*
32. Roy PC, Boddeti VN (2019). Mitigating information leakage in image representations: a maximum entropy approach. In: *IEEE conference on computer vision and pattern recognition (CVPR)*, pp 2586–2594
33. Zheng L, Zhang Y, LL Thing V (2019) A survey on image tampering and its detection in real-world photos. *J Vis Commun Image Represent* 58:380–399
34. AP da Costa K, Papa JP, Passos LA, Colombo D, Del Ser J, Muhammad K, de Albuquerque VHC (2020) A critical literature survey and prospects on tampering and anomaly detection in image data. *Appl Soft Comput*
35. Meena KB, Tyagi V (2019) Image forgery detection: survey and future directions. In: *Data, engineering and applications*, pp 163–194. Springer
36. Nowroozi E, Dehghantanha A, Parizi RM, Choo KR (2020) A survey of machine learning techniques in adversarial image forensics. *Comput Secur*
37. Bulusu S, Kailkhura B, Li B, Varshney P, Song D (2020) Anomalous example detection in deep learning: a survey. *IEEE Access* 8:132330–132347
38. Wang X, Li J, Kuang X, Tan Y, Li J (2019) The security of machine learning in an adversarial setting: a survey. *J Parallel Distrib Comput* 130:12–23
39. Chakraborty A, Alam M, Dey V, Chattopadhyay A, Mukhopadhyay D (2018) Adversarial attacks and defences: a survey. *CoRR*, abs/1810.00069
40. Agarwal A, Singh R, Vatsa M, Ratha NK (2020) Image transformation based defense against adversarial perturbation on deep learning models. *IEEE transactions on dependable and secure computing (TDSC)*
41. Baker S, Kanade T (2000) Hallucinating faces. In: *IEEE international conference on automatic face and gesture recognition (FG)*, pp 83–88
42. Grm K, Pernuš M, Cluzel L, Scheirer WJ, Dobrsek S, Štruc V (2019) Face hallucination revisited: an exploratory study on dataset bias. In: *IEEE conference on computer vision and pattern recognition workshops (CVPR)*
43. Grm K, Scheirer WJ, Štruc V (2020) Face hallucination using cascaded super-resolution and identity priors. *IEEE transactions on image processing (TIP)*, 29:2150–2165
44. Othman A, Ross A (2014) Privacy of facial soft biometrics: suppressing gender but retaining identity. In: *European conference on computer vision (ECCV)*, pp 682–696. Springer
45. Terhöst P, Damer N, Kirchbuchner N, Kuijper A (2019) Suppressing gender and age in face templates using incremental variable elimination. In: *International conference on biometrics (ICB)*, pp 4–7
46. Johnson D, Sinanović S (2001) Symmetrizing the Kullback-leibler distance. *IEEE Transactions on information theory (IT)*
47. Mustafa A, Khan SH, Hayat M, Shen J, Shao L (2019) Image super-resolution as a defense against adversarial attacks. *IEEE Trans Image Process (TIP)*, 29:1711–1724
48. Huang GB, Ramesh M, Berg T, Learned-Miller E (2007) Labeled faces in the wild: a database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst
49. Milborrow S, Morkel J, Nicolls F (2010) The MUCT landmarked face database. Pattern recognition association of South Africa. <http://www.milbo.org/muct>
50. Eidinger E, Enbar R, Hassner T (2014) Age and gender estimation of unfiltered faces. *IEEE Trans Inform Forensics Secur (TIFS)*, 9(12):2170–2179
51. Chatzikyriakidis E, Papaioannidis C, Pitas I (2019) Adversarial face de-identification. In: *IEEE international conference on image processing (ICIP)*, pp 684–688

52. Simonyan K, Zisserman A (2015) Very deep convolutional networks for large-scale image recognition. CoRR, abs/1409.1556
53. Cao Q, Shen L, Xie W, Parkhi OM, Zisserman A (2010) VGGFACE2: a dataset for recognising faces across pose and age. In: IEEE international conference on automatic face and gesture recognition (FG), pp 67–74
54. Mirjalili V, Raschka S, Ross A (2018) Gender privacy: an ensemble of semi adversarial networks for confounding arbitrary gender classifiers. IEEE international conference on biometrics theory, applications and systems (BTAS), pp 1–10
55. Yi D, Lei Z, Liao S, Li SZ (2014) Learning face representation from scratch. CoRR abs/1411.7923, 2014
56. Meden B, Emeršič Ž, Štruc V, Peer P (2018) k-Same-Net: k-anonymity with generative deep neural networks for face deidentification. Entropy 20(1):1–24

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.