Yunqian Wen · Bo Liu · Li Song ·
Jingyi Cao · Rong Xie

Face De-identification: Safeguarding Identities in the Digital Era

Springer

Face De-identification: Safeguarding Identities in the Digital Era

Yunqian Wen • Bo Liu • Li Song • Jingyi Cao •
Rong Xie

Face De-identification: Safeguarding Identities in the Digital Era

Springer

Yunqian Wen
Department of Electronic Engineering
Shanghai Jiao Tong University
Shanghai, China

Li Song
Department of Electronic Engineering
Shanghai Jiao Tong University
Shanghai, China

Rong Xie
Department of Electronic Engineering
Shanghai Jiao Tong University
Shanghai, China

Bo Liu
School of Computer Science, University of Technology Sydney
Ultimo, NSW, Australia

Jingyi Cao
Department of Electronic Engineering
Shanghai Jiao Tong University
Shanghai, China

ISBN 978-3-031-58221-9 ISBN 978-3-031-58222-6 (eBook)
https://doi.org/10.1007/978-3-031-58222-6

“In recognition of those committed to safeguarding identities and advancing privacy in the digital realm. Your dedication to the ethical use of technology shapes a future where innovation coexists harmoniously with personal privacy.”

Preface

Welcome to "Face De-identification: Safeguarding Identities in the Digital Era." As the author/editor of this book, I am honored to present this comprehensive exploration into the intricate realm of safeguarding identities in an increasingly digital landscape.

The idea for this book stemmed from a deep-rooted concern for privacy and security in today's technologically advanced world. The scope of this work encompasses an extensive study of face de-identification techniques, aiming to address the critical challenges faced in protecting identities amid the pervasive use of facial recognition technologies.

Our intent with this book is to offer a thorough examination of various face de-identification methodologies, elucidating their intricacies, strengths, and limitations. Through a structured approach, we have endeavored to present an array of techniques, from obfuscation-based methods to advanced deep generative models, catering to a diverse audience interested in understanding the multifaceted aspects of preserving privacy in digital spaces.

This book is designed for scholars, researchers, practitioners, policymakers, and individuals curious about the intersection of technology and privacy. It serves as a resource for academics delving into the complexities of identity protection, professionals implementing privacy measures, and enthusiasts seeking a deeper understanding of face de-identification in an evolving digital world.

Sydney, NSW, Australia — Bo Liu
November, 2023

Acknowledgments

This work was supported by the Fundamental Research Funds for the Central Universities, the MoE-China Mobile Research Fund Project (MCM20180702) and the National Key R&D Project of China (2019YFB1802701).

About the Book

Our combined team from University of Technology Sydney (UTS) and Shanghai Jiao Tong University (SJTU) started to work on the topic of face de-identification from 2020. Our findings in state-of-the-art de-identification technology have been invaluable, making the insights and perspectives highly commendable and respected.

In this compelling work, the reader is presented with an insightful journey into the world of face de-identification. As the team's engaging narrative unfolds, you will be guided through the intricate landscape of safeguarding identities in the digital era.

As an expert in privacy protection, I have witnessed the evolution and impact of technologies on our daily lives, especially with respect to privacy and security concerns. The exploration of face recognition and de-identification techniques in this book is timely and essential in our increasingly interconnected world.

This book introduces a comprehensive exploration of face de-identification techniques, shedding light on the complexities and challenges faced in this field. The innovative strategies and ethical considerations presented here mark a significant step forward in the ongoing dialogue on privacy and identity protection. I am confident that this work will contribute significantly to the discourse on privacy and technology, fostering deeper insights and inspiring further advancements in this crucial area.

I commend all my co-authors for their dedication and expertise in compiling this significant contribution. It is my privilege to introduce this impactful work to readers and commend its relevance, timeliness, and scholarly merit.

Contents

Acronyms

ϵ	Privacy Budget
2D	Two-Dimensional
3D	Three-Dimensional
3DMM	3D Morphable Model
A^3GAN	Attribute-aware Anonymization Network
AAD	Adaptive Attentional Denormalization
AAM	Active Appearance Model
AE	Auto Encoder
AINet	Attribute-aware Injective Network
AU	Action Unit
BFR	Blind Face Restoration
BIPA	Biometric Information Privacy Act
CA	Coded Aperture
CCPA	California Consumer Privacy Act
cGAN	conditional Generative Adversarial Network
CNN	Convolutional Neural Networks
CS-SFT	Channel-Split Spatial Feature Transform
DGN	Deep Generative Network
DNN	Deep Neural Network
DP	Differential Privacy
DRRDN	Deep Robust Representation Disentanglement Network
FACS	Facial Action Coding System
FATM	Facial Attribute Transfer Model
FID	Fréchet Inception Distance
FIP	Facial Identity-Preserving
GAN	Generative Adversarial Network
GDPR	General Data Protection Regulation
HiSD	Hierarchical Style Disentanglement
ISR	Inverse Super Resolution
JPEG	Joint Photographic Experts Group
LDP	Local Differential Privacy

LPIPS	Learned Perceptual Image Patch Similarity
MAE	Mean Absolute Error
MfM	Multi-factor Modifier
MMDA	Multimodal Discriminant Analysis
NeRF	Neural Radiance Field
OPOM	One Person One Mask
PATE	Private Aggregation of Teacher Ensembles
PCA	Principal Component Analysis
PCC	Pearson's Correlation Coefficient
PDPA	Personal Data Protection Act
PI	Perceptual Indistinguishability
PIPEDA	Personal Information Protection and Electronic Documents Act
PPAS	Privacy-Preserving Attribute Selection
PSNR	Peak Signal-to-Noise Ratio
R^2VAE	Replacing and Restoring Variational Autoencoder
RMSE	Root-Mean-Square Error
SOTA	State-of-the-Art
SSIM	Structural Similarity Index Measure
SVD	Singular Value Decomposition
VAE	Variational Auto-Encoder
WGAN	Wasserstein Generative Adversarial Network

Part I
Introduction

Chapter 1
Introduction

1.1 Background and Motivation

In recent years, the world has borne witness to a rapid surge in artificial intelligence technologies, particularly those rooted in deep learning, alongside the widespread proliferation of face recognition applications. This technological renaissance, however, brings with it a pressing concern—privacy [1–4]. Amidst these groundbreaking advancements, faces stand out as one of the most sensitive forms of biological information, intimately connected to personal identity. The essence of face recognition lies in its biometric authentication, a characteristic that is both unique and irrevocable. Yet, the consequences of this technology extend far beyond mere identity verification. On the one hand, when harnessed for cross-referencing with other databases, it unveils a wealth of an individual's sensitive information. A landmark study by Acquisti et al. [5] underscored how faces can serve as the link connecting diverse databases, revealing trails associated with various personas and ultimately undermining privacy. On the other hand, after confirming the identity of a face through face recognition technologies, advanced visual analysis and understanding tools can infer a large amount of sensitive privacy information from the corresponding visual face. For instance, occupation [6] and health status [7]. This poses a serious threat to the security of personal information.

In light of these growing privacy concerns, the field of face de-identification has emerged as a vital research domain within the realms of security and privacy. Face de-identification, a process that conceals facial features while preserving utility for identity-unrelated applications, has found applications in a wide range of scenarios, from anonymizing faces in media interviews and video surveillance [6] to safeguarding privacy in medical research [7], and beyond [8, 9].

The ubiquity of image acquisition in our daily lives—be it sharing personal images on social media, online learning with cameras, or public safety surveillance—renders the need for enhanced privacy protection all the more critical. Existing privacy safeguards often prove inadequate, allowing third parties to collect

Y. Wen et al., *Face De-identification: Safeguarding Identities in the Digital Era*,
https://doi.org/10.1007/978-3-031-58222-6_1

human facial images without consent for large-scale data analysis or questionable applications.

Prominent social media platforms like Google, Facebook, and Shutterfly have faced scrutiny for compromising the privacy of millions by inadvertently leaking private photos to commercial entities, thus embroiling themselves in biometric privacy disputes. Conversely, the need for extensive public facial image datasets to fuel the development of cutting-edge deep learning models has led to the creation of invaluable resources. Yet, these repositories carry inherent privacy risks, resulting in increased restrictions on data sharing. Notably, datasets such as Microsoft's MS-Celeb-1M, Duke's MTMC, and Stanford's Brainwash were, at various times, withdrawn from public access due to privacy concerns.

The growing spotlight on privacy issues has prompted the enactment of stringent laws and regulations, notably the General Data Protection Regulation (GDPR) [10, 11], which prohibits companies from collecting, sharing, or analyzing user data without informed consent. Within the GDPR framework, privacy information encompasses "personal data related to an identified or identifiable natural person," underscoring the paramount importance of protecting personal identity, particularly in the context of facial image data.

This book, "Face De-identification: Safeguarding Identities in the Digital Era," endeavors to explore the multifaceted landscape of face de-identification. It delves into a wide array of methods and strategies aimed at preserving facial sensitive information, notably identity, while retaining utility for applications unrelated to identity. Through a comprehensive examination of this crucial field, we seek to provide both practitioners and researchers with the knowledge and tools necessary to navigate the intricate intersection of technology, privacy, and identity protection.

1.2 Face Recognition and Face De-identification

From the background and motivation, it can be seen that face de-identification is a benign technology born to stop face recognition from invading personal privacy, and the two are in a state of confrontation with each other. In order to design excellent face de-identification technology, a thorough understanding of face recognition technology is a necessary condition. Therefore, next we will introduce face recognition and face de-identification separately.

1.2.1 Face Recognition

Face recognition is a biometric technology that automatically recognizes people's facial features including statistics and geometric features, which is one of the most important applications of image analysis and understanding. Face recognition tasks can be further divided into binary classification and multiclassification. The binary

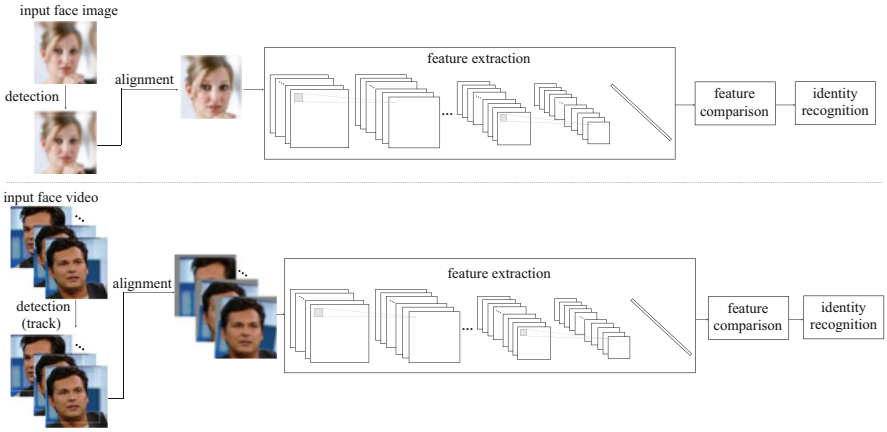

Fig. 1.1 The face recognition process. First, the input image or video is detected and possibly tracks (just for video) to localize the faces. Second, the detected faces are aligned to normalized canonical coordinates. Third, deep facial features are extracted by various methods. After well-designed feature comparison, the identity of the input face data is finally recognized

classification task is also called face verification, which is used to compare whether two images have the same identity. The multiclassification task is also called face retrieval, such as searching for a face with a specific identity in a database of many faces. The widely known face recognition is the abbreviation for identity recognition and verification based on optical facial images. The face recognition process can be simply summarized as using a computer to analyze a face video or image. Firstly, it detects and possibly tracks (just for videos) the faces, so as to localize them. Secondly, it aligns the faces to normalized canonical coordinates. Thirdly, it extracts effective facial features. Finally, it determines the identity of the face object through a comparison of the above-mentioned features. The whole process is shown in Fig. 1.1.

The research on face recognition can be traced back to the late 1960s. The main idea is to design feature extractors and then use machine learning algorithms for classification. Traditional methods rely on hand-made features, such as edge texture description, and combine with machine learning techniques such as principal component analysis, linear discriminant analysis, and support vector machines. The early methods based on geometric features focused on extracting contours and geometric relationships of face components and using the geometric descriptions of shapes and structural relationships as features to construct several feature vectors, including the distance, curvature, and angle between two specified facial keypoints. The advantages are fast recognition speed and low requirements of memory, while the disadvantages are that geometric features can only describe the basic facial information, ignore local subtle features, and result in the loss of local information. The current feature point detection technology is far from meeting the requirements in terms of accuracy.

After introducing deep learning techniques into the field, the approaches have been transferred to extract features with neural networks, which has greatly improved the accuracy and robustness. The deep learning models can be trained by a large amount of data to learn the representation of various variability such as lighting conditions, postures, facial expressions, and so on.

Today, face recognition technology has been widely used in our daily life. Face verification can be treated as a new way of identity confirmation for fast face comparison, mobile payment authentication, security identity verification, etc. Face retrieval can be applied to investigate suspects, complete search of missing persons' databases, and repeated investigation of multiple certificates for one person. At present, the face recognition model can achieve satisfactory accuracy on a specific dataset, but the influence of illumination and posture is still the main challenge. In addition, cross-racial and cross-age recognition problems are also worth studying.

1.2.2 Face De-identification

Due to potential privacy issues, the application of face recognition technology is currently under controversy, and the face privacy protection task is receiving more and more attention. Face de-identification, the main content of this book, is an innovative technical idea to solve the dilemma. There is no consistent definition of de-identification in the existing literature. Ribaric et al. [12] defined de-identification in multimedia content as *"the process of concealing or removing personal identifiers, or replacing them with surrogate personal identifiers in multimedia content."* During this process, other facial features that are not related to identity should remain unchanged, such as expression, posture, and background. After this process, the de-identified face will be judged by the face recognition technology as no longer the same identity as the original face. At the same time, the identity-protected face is expected to retain as much similarity to the original image as possible for normal viewing and sharing and can still be analyzed and processed by other identity-agnostic computer vision methods, such as face detection, motion monitoring, and emotion recognition. Additionally, better image quality and visual effects are also preferred.

With face de-identification technologies, visual service providers can use face visual data to carry out legitimate scientific research, business analysis, security monitoring, social sharing, and other activities; ordinary individuals can enjoy the convenience of visual technology without worrying about their other biometric information due to personal identity associated with the disclosure. It effectively alleviates the concerns about personal privacy and security in today's society. To sum up, providing identity protection for facial visual data is the trend of our time, which has great social significance and practical value.

It is recognized that the main purpose of face de-identification is to conceal the identity information of a face. Images and videos are the two main visual data of human faces, and they are also the focus of face de-identification research.