

# Understanding Person Identification Through Gait

Simon Hanisch<sup>1</sup> Evelyn Muschter<sup>2</sup> Admantini Hatzipanayioti<sup>2</sup>  
Shu-Chen Li<sup>2</sup> Thorsten Strufe<sup>1,3</sup>

<sup>1</sup>Chair of Privacy and Security, Centre for Tactile Internet with Human-in-the-Loop (CeTI), TU Dresden

<sup>2</sup>Chair of Lifespan Developmental Neuroscience, CeTI, TU Dresden

<sup>3</sup>Karlsruhe Institute of Technology (KIT), Germany

<sup>1,2</sup>{firstname}.{lastname}@tu-dresden.de, <sup>3</sup>thorsten.strufe@kit.edu

## ABSTRACT

Gait recognition is the process of identifying humans from their bipedal locomotion such as walking or running. As such, gait data is privacy sensitive information and should be anonymized where possible. With the rise of higher quality gait recording techniques, such as depth cameras or motion capture suits, an increasing amount of detailed gait data is captured and processed. The introduction and rise of the Metaverse is an example of a potentially popular application scenario in which the gait of users is transferred onto digital avatars. As a first step towards developing effective anonymization techniques for high-quality gait data, we study different aspects of movement data to quantify their contribution to gait recognition. We first extract categories of features from the literature on human gait perception and then design experiments for each category to assess how much the information they contain contributes to recognition success. We evaluated the utility of gait perturbation by means of naturalness ratings in a user study. Our results show that gait anonymization will be challenging, as the data is highly redundant and inter-dependent.

## KEYWORDS

biological motion, privacy, gait, identification, recognition, anonymization

## 1 INTRODUCTION

Human gait is a biometric trait that can be used to identify persons, infer private information such as sex [45] or age [54], and is used in the diagnosis of medical conditions. When compared to person identification via faces, gait has advantages as it can be done from distances at which the face is not yet recognizable or occluded by objects such as face masks. It is believed that distinguishing individuals from afar was an important human survival mechanism in the past, as it allowed to recognize if an individual was a friend or foe before the person was close enough to be a potential threat [53].

In today's world, we no longer need to rely as much on gait recognition for our own safety. However, the ease of gait recognition, given near ubiquitous means of capturing and recording people, creates novel threats to privacy. Examples from China<sup>1</sup> show that

<sup>1</sup><https://tinyurl.com/5ya4cwdd> [apnews], accessed: 17.08.2022

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.  
Proceedings on Privacy Enhancing Technologies YYYY(X), 1–14  
© 2023 Copyright held by the owner/author(s).  
<https://doi.org/XXXXXXX.XXXXXXX>



gait is very much suited to be used for surveillance purposes alongside face recognition. As for the future, human gait will increasingly be captured and processed (cf. Fig. 1 for recent examples), for example, to realize visions like the Metaverse [14]. The Metaverse is an immersive virtual world for which human behavior, including human gait, is captured and then transferred onto digital avatars. Another example is safe human-robot collaboration [38] in which humans have to be closely tracked in order to avoid collisions with robots. The technology used for these scenarios captures the human motion either via vision-based approaches using cameras (e.g., Kinect<sup>2</sup>) or wearable systems based on inertial measurement units (IMUs) (e.g., SmartSuit Pro<sup>3</sup>).



**Figure 1: Variety of recent motion capturing devices recording gait of individuals in public and private spaces from video or inertial measurement units (cf. OpenPose, VFXVoice, Virtuix).**

In order to preserve the privacy of individuals, anonymization methods are required to remove or hide sensitive information of their gait. We work towards this goal by investigating which features of human gait make individuals identifiable or allow the inference of personal attributes like sex. As a starting point for finding categories of features for our computational experiments, we look at the literature on human gait perception. This line of research has long been a topic of cognitive science and investigates how humans identify other humans by their gait (e.g., [25, 49, 11]). We would like to emphasize that we are not interested in designing a novel attack, and we are less interested in investigating the robustness of specific anonymization schemes that are hard to interpret. Instead, we are interested in the question of which

<sup>2</sup><https://azure.microsoft.com/en-us/services/kinect-dk/>

<sup>3</sup><https://www.rokoko.com/products/smartsuit-pro>

features in precise gait data yield identity or attribute disclosure of the individuals, if they coincide with those known from psychological research on gait and person perception—and to which extent they are inter-dependent and hence cannot independently be suppressed/perturbed for anonymization.

For a systematic analysis we use machine learning (ML) to get an estimate of how much identifying information exists in gait data. We then design perturbations for each of the feature categories that remove specific features in the gait data, to then measure how much the recognition performance drops. Where possible, we try to manipulate a feature alone, so that we can estimate how much identifying information the feature contains, as well as how much it shares with other features. To establish how much utility is retained after anonymization, we performed a user study in which the participants rated the naturalness of the resulting gaits. Key contributions of our work are as follows:

- human gait feature categorization extracted from the literature of cognitive science;
- systematic study of feature contribution for gait recognition for both identification and sex classification;
- simple gait perturbation techniques;
- utility evaluation via perceived naturalness of the anonymized gait.

In the following, we start by first explaining the system model we follow and then review related work from the fields of computer science and cognitive science in Section 2. This includes the relevant literature in human gait recognition and person perception in the field of cognitive science. In Section 3 we describe our general methodology and the design and implementation of our computational experiments. We present the experimental results in Section 4. Finally, Sections 5 and 6 offer a discussion and conclusion.

## 2 BACKGROUND AND RELATED WORK

In the following section we provide background on current motion capturing including potential threats, the human and automatic recognition tasks based on gait, and the state of the art with regards to anonymization and explainability-based analyses of identifying features in gait.

### 2.1 System Model and Threats

Our main interest is to identify those features in gait data that carry the most information for the identification of individuals, or to infer their attributes, like sex. Our interest is based on the observation that motion is increasingly captured and processed for various purposes. Consider a user who is using a motion capture system to transfer their motions onto a digital representation (e.g., an avatar or digital twin). Recent examples include various console games, but also the current vision of the Metaverse. These examples require the digital motion to mimic the motions of the user in a natural looking way, resembling them as closely as possible.

The movements of the user are captured locally in real-time, based on either visual recording or wearables with embedded sensors. It is then locally pre-processed, for example to extract motions from the video or sensory data, and then transmitted to the service

provider, who analyses and distributes it to their customers and other users.

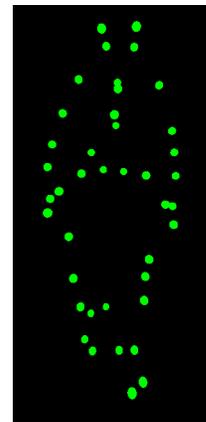
Recognizing or identifying a person in such services clearly has social relevance [12, 40] and clinical power for certain disorder diagnostics (e.g., stroke, Parkinson’s disease) [13, 54]. This immediately implies that there are privacy issues to be considered, both in terms of recognition, for instance for surveillance, but also in terms of attribute inference.

The *service provider* on the one hand will likely be in the position of identifying and linking subsequent observations of an individual based on their account. However, the processed motion data may yield inferences of attributes that are private and not relevant to the service, and the possibility of linking several independent accounts of the same individual. On the other hand, *third parties* e.g., other users or platform and advertising customers of the provider, will be able to train recognition models of users in order to re-identify (for instance in auxiliary video data), or also learn sensitive attributes that are not shared voluntarily by the individual.

Hence, we aim to understand which features of the motion provide high predictive power for identity or attribute disclosure. This would facilitate the development of local privacy-enhancing technologies for the pre-processing of data prior to their transfer. Understanding, which features carry information required for the motion to be perceived as natural will help strike useful privacy-utility tradeoffs.

### 2.2 Motion Analysis

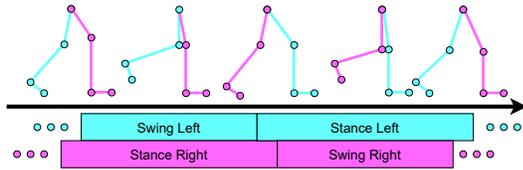
Humans can recognize and identify biological traits visually through the use of static information such as shape or other cues. Biological motion is one additional important factor. Human newborns, infant monkeys, and even freshly hatched, visually naïve chicks show a preference for motion cues that move in motion patterns suggesting animacy [31].



**Figure 2: Example of a human walker represented as a point-light display (PLD).**

An established and reliable method of investigating biological motion processing without other potential sensory information or cues (e.g., color, texture, or form-based features such as facial configurations, hairstyle, or clothing) is the use of point-light displays

(PLDs) [25] (see Fig. 2). This method of using impoverished moving dot displays helps to isolate motion information from other cues. Thereby, a small number of dots represent the head and major joints of a human body in various scenarios such as social interactions [7, 33] or—as the focus of this work—during gait [49]. Indeed, many identifying features of a human walker can be recognized from such PLDs.



**Figure 3: A schematic representation of the gait cycle.**

*Gait analysis* is the study of human locomotion (walking and running) and defines walking as a series of gait cycles. A *gait cycle (stride)* is the period when one foot contacts the ground to when that same foot contacts the ground again (see Fig. 3). Each gait cycle has two phases: the stance phase, when the foot is in contact with the ground; and the swing phase when the foot is not in contact with the ground. In vision-based gait analysis such as PLDs, kinematic data such as position and velocity are captured. These are used to relate motion parameters such as joint angles, joint velocity, and center of mass with qualitative gait parameters—the general prototypical bipedal motion characteristics such as step length, walking speed, pace and rhythm of steps, stance, and swing times [13] as well as arm swing, vertical head movement, pelvic rotation, and the extension and flexion of limbs and shoulders. As walking consists of a series of multiple gait cycles, gait data typically also contains fluctuations. These are small variations such as asymmetry and variability in step and stance time, step velocity, or step length [13]. These gait parameters can subsequently be used to extract statistical features (e.g., mean, standard deviation, skewness) for gait analysis.

### 2.3 Human Gait Perception

Human observers have no trouble making sense of the very limited information presented through PLDs’s disconnected dots, representing actions such as the specific categorical biological motion content [53, 40] of human gait (walking or running). Research has suggested that humans are especially tuned to recognizing conspecifics and this preference is likely to already emerge in the visual system. Psychophysical and neuroscientific studies have shown that at least two processes play a role in person perception (here defined as the recognition of human bodies and their biometric features based on vision).

On the one hand, form-from-motion cues [46, 53] are cues that are rooted in basic perceptual abilities to see structure from motion. That is, the shape and form of an object or person are revealed more clearly through motion. The human visual system benefits from the motion direction information in order to extrapolate the overall shape of an object or person. These cues provide time-invariant information about body form by enhancing the shape presentation

of a person [53] and are susceptible to violations of the hierarchical body form structure [10, 42] as well as to inversion effects [18]. That is, inverting a body in the image plane (i.e., placing it upside-down) results in perceptual impairments. Previous research has proposed that inversion is deleterious to normal human whole-body perception, causing observers instead to rely on local part-based visual features [18]. Evidence for first order configural processing has shown that visual perception of bodies is mediated by spatial configurations of body parts, such as the general body layout (e.g., legs attached to the hip, arms attached to the shoulders), and thus providing intact spatial configurations of bodies [10].

On the other hand, dynamic identity signature [46, 53], describes the idiosyncratic motion pattern of an individual. These features describe the change over time during a walking cycle and rely on nuanced, person-specific motion variations (e.g., the way Charlie Chaplin walks). Furthermore, research has provided evidence for a two-stream processing of biological motion perception in the brain. That is, biological motion perception relies on—both dynamic and static features through—motion processing in the dorsal pathway (i.e., area V5 of visual cortex in the brain) in combination with bodily form and appearance information in the ventral pathway in the brain (see Peng et al. [42] for further details on visual recognition of biological motion). In addition to action recognition, human observers are able to identify soft biometric features of actors in PLDs, including sex [37, 18, 28], age [54, 37], weight [39], height [39], handedness [30], in addition to attractiveness [37, 39], identity [12], emotions [37, 26] and causal intentions [33]. Specifically, Kozlowski and Cutting [28] showed that the biomechanical factor center of moment, which is derived from the relative movement of both – the shoulders and hips, plays a crucial role in sex perception in PLDs.

Finally, person recognition depends on familiarity and might take time for the human observer to learn, but is useful for recognizing a familiar person from a distance [12, 49]. Studies have shown that human observers are more sensitive to PLDs of themselves and friends [12, 32], or could learn to identify a small number of individuals based on their motion [49]. Guided by these insights, we aim to investigate if the removal of the certain features will reduce recognition rates, and to which extent. Specifically, we focus on the features that are easy to extract from existing data sets. Namely, macro and micro features (i.e., statistical features, see Sect. 2.2), perturbations of intact bodies in natural spatial configurations, as well as dynamic (i.e., temporal information) and static (i.e., structural) features. Section 3.3 will describe the specific features used in the present study in detail.

### 2.4 Automatic Gait Analysis

Current human movement analyses are based on biometric measurements and motions. They are captured vision-based, or using wearables with integrated inertial measurement units (IMUs). A gait cycle is thereby composed of a chain of individual 2D (video) or 3D (optical marker/IMU tracking) samples at each given time point (pose).

Gait recognition using machine learning models is most commonly based on video data [51]. Video, providing rich information about subjects, facilitates high recognition rates and hence is frequently used for surveillance purposes [5].

Also explicit motion capturing frequently uses video: High-quality vision-based motion capturing uses specialized cameras to track reflective markers on the subject’s body. The position of these markers is later reconstructed into 3D position time series, converted into joint angles as a function of time, and subsequently analyzed according to specific research or clinical needs. However, recently, approaches of using a single commodity camera in combination with keypoint detection algorithms and neural networks (e.g., Open Pose or DeepLabCut, cf. Fig. 1) have generated convincing results [27, 36, 34].

Gait recognition is also possible based on motion capturing data [4]. Indeed, even simple kinematic features obtained from IMU systems (e.g., position, velocity, and acceleration-based features) or kinetic data from force plates and electromyography (e.g., ground or muscle force parameter) have been shown to yield high recognition rates (see Connor and Ross [11] for an overview).

Anonymizing individuals in video surveillance footage for multiple moving object detection and tracking algorithm (e.g., human action tracking) by representing their bodies as simplified objects such as PLDs thus cannot protect their identities. Further, gait can also be used to infer personal attributes like sex [45] and age [54, 15]. Being interested in those gait features that carry information for identification and attribute disclosure of individuals, in the present work we rely on marker-based motion capture data as it is considered the gold standard in the field.

### 2.5 Existing Anonymization Approaches

As protection against gait recognition, gait anonymization recently became an active field of research. However, in the present paper we are not interested in evaluating their efficacy, but rather want to establish an initial understanding as to which part of gait data carries identifying information. Yet, we use the state of the art on anonymization for guidance.

For example, Ivasic-Kos et al. [24] anonymized activity videos by blurring the data. This clearly retains the gait information in the data and does not facilitate insights into its influence on identification. More sophisticated methods use neural network approaches to perturb the video recordings of gait and then generate a new natural-looking gait sequence [48, 21]. While manipulating, and hence implicitly providing indicators about the influence of different properties in gait, neither approach investigates the actual influence but rather provides ad-hoc anonymization. For accelerometer-based data, it was proposed to add noise directly to signals collected via smartphones [35]. This perturbation may have some effect given singular smartphone-based measurements. However, it can be removed from the data, as Wang et al. [52] have shown for trajectories, and it does not provide insights as to which details of the data are identifying the users.

### 2.6 Investigations using Explainability

The field of explainable machine learning [44, 2] focuses on the interpretability of learned machine learning models. It helps to determine those features that are important for the classification, for instance of gait recognition systems.

A common approach is layer-wise relevance propagation (LRP) [3]. It is applied to neural networks (NN) to find the most relevant

features in the input data by tracing the classification back through the layers of the NN. Another approach is to perturb areas in the input data in order to find the area that results in the largest classification performance reduction [16]. The corresponding insights are limited to implicit indicators to regions of the input data that influence the classification accuracy. They do not provide insight as to the explicit semantic feature of the data, in our case—the different characteristic features of gait, that contribute to identification of the individuals.

Approaches for explainable machine learning have also been used to analyze gait patterns for clinical analyses [47]. Horst et al. [23] used LRP to study which part of the gait cycle is relevant to a non-linear machine learning model to recognize an individual. Furthermore, Connor and Ross [11] investigated which features in gait data make people identifiable. These approaches test trained models for effects of the removal of features, without retraining them. The question remains, whether other features were considered unnecessary and thus ignored by the initially trained model, as all of their information is also contained in the features used by the model.

Alvarez Melis and Jaakkola [1] proposed faithfulness as an important metric for evaluating explainable machine learning. Faithfulness is obtained by removing/perturbing the feature and then measuring the drop in classification performance. While their investigation does not shed light on the contribution of explicit gait features, it does provide instructive ideas for our analysis.

It can be concluded that several studies have investigated gait recognition by both humans and machines, as well as implicit approaches to anonymize gait. However, so far the question of which explicit features of gait have predictive power for identification tasks—and hence have to be perturbed or removed for anonymization—has not been investigated systematically.

## 3 METHODS

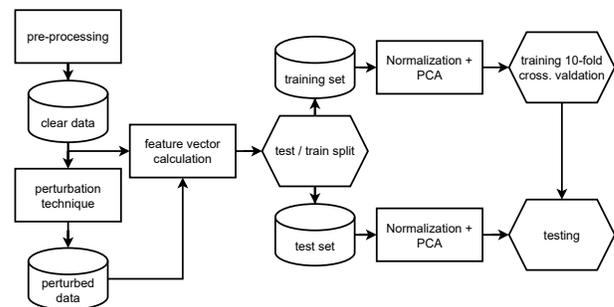


Figure 4: The full data processing pipeline.

The question we sought to answer is how much specific features in the data contribute to the overall gait recognition performance of identity and sex using machine learning. Our overall approach was to first train & test a gait recognition system for each of the recognition goals on clear data to obtain baseline accuracy. Next we obfuscated a feature at a time in the data by either perturbing or removing it to investigate its impact on anonymization. We then repeated the training & testing process and report the resulting

recognition performance. The difference in recognition between baseline and perturbed data gives us the unique amount a feature contributes to the overall recognition performance. Further, we also measured each feature independently from the other features. However, this is only possible for features we can remove from the data and not for features we only perturb. The full process is show in Fig. 4.

In the following, we provide details about the data set, applied feature perturbations, the implementation of the recognition system, and utility evaluation.

### 3.1 Data Set

As our main goal was to understand the important features of human gait, we chose the highest quality of gait data and used optical 3D marker-based motion capture data for our experiments. This data is considered the gold standard for motion capturing and is recorded using multiple infrared cameras which capture markers on the anatomic landmarks of participants. The benefit of the 3D representation is that there is no dependency on the recording angle like in video recordings. The data consists of multiple samples per participant which are a time series of poses. Each pose contains the 3-dimensional coordinates of each marker (placed on the participant) at a given point in time (i.e., PLDs). The data is also more appropriate for our purpose, as we focus on gait features in the absence of potential additional information (e.g. video recordings). Furthermore, one can generate PLDs from the data, using the tracked marker positions.

We used the open-source data set by Horst et al. [22, 23] which consists of full-body kinematic and kinetic data of 57 individuals (29 female, 28 male;  $23.1 \pm 2.7$  years;  $1.74 \pm 0.10$ m;  $67.9 \pm 11.3$ kg). An optical motion capture system and a full body marker set (62 markers corresponding to anatomical landmarks), as well as two force plates, recorded self-paced walking trials at 250Hz (motion capture) and 1000Hz (force plates). For each participant 20 samples containing a full gait cycle have been recorded (for further details on the data acquisition protocol see Horst et al. [23]).

### 3.2 Data Pre-processing

Following the methodology of Horst et al. [23], we trimmed the gait samples to contain only a single stride by using the kinetic force signals of the force plates, using a ground force threshold of 20N. This way all samples are aligned and start at the same point in the gait cycle. The data was then normalized, in order to obtain an equal number of poses for each individual, by resampling each sample to 100 frames. Each frame represents one discrete pose of the individual while walking, the 100 poses then constitute one stride.

### 3.3 Retained and Masked Features

We based our feature categories on previous work in gait analysis and human perception as described in Sections 2.2 and 2.3. The category name always gives the kind of feature we sought to retain, while the perturbation techniques employed are aimed at removing the other features from the data. For each technique, we strove to design an inverse perturbation technique that only removes the

specific feature, while keeping all the others (micro vs. macro, dynamic vs. static). This way we sought to understand how much each feature contributes to the overall recognition rate and if it contains information that is unique to this feature. Since there is interdependence between features, some of the features are partially overlapping for example the walking frequency is dependent on the walking speed and the length of the legs. Table 1 gives a brief overview, while the used and obfuscated features are described in detail in the following.

Our **macro** features describe the general characteristics of the walker, such as walking speed, general movement trajectories, walking amplitude, the most significant parts of the walker positions, and overall body parts. Its counterparts are the **micro** features which contain the small variations of the trajectories that remain when the overall trajectories are removed, the walker without its walking speed and step length equalized over all walkers, the least significant parts of the walker positions, and individual body parts. Besides macro and micro, we also investigated the dynamic parts of the gait motion. For this we have two contrary feature categories static and dynamic. The **static** features contain the time-invariant features, such as the average pose of the walker, or the first pose of the walker. The **dynamic** features contain the features describing the motion of the walker, including the differences between the recorded poses, and walker where the static frame (body proportions) has been removed. The following section describes the used perturbation techniques for each feature category. The parameter values have been chosen to match the used data set. In the end, we briefly detail how we combined the perturbation techniques.

**Table 1: used and obfuscated features**

	Macro	Micro	Static	Dynamic
<i>definition</i>	<i>Step length, walking speed, cadence</i>	<i>asymmetry and variability in the macro features</i>	<i>Shape and general body layout</i>	<i>Time course of changes</i>
Perturbation 1	Remove variations	Remove trajectories	Static pose	Motion extraction
Perturbation 2		Amplitude/frequency equalization	Resampling	Normalization
Perturbation 3	Coarsening macro	Coarsening micro		
Perturbation 4	Remove body parts	Keep body parts		

We provide a sample video rendering<sup>4</sup> of all perturbation methods alongside this paper.

**3.3.1 Macro Features.** The macro features keep the overall characteristics of the walker and remove its smaller variations from the data. We used three perturbation techniques for this: remove variations, coarsening macro, and remove body parts.

**Remove variations:** In order to extract the ideal trajectory from the gait data we removed the small variations that deviate from the ideal trajectory. The ideal trajectory is here calculated by two different methods: either using a moving window on the marker poses and then calculating a rolling average, or an interpolation. The difference between the two is that the rolling average takes all poses in the window to calculate an average, while the interpolation only uses the poses at the edge of the moving window. The moving window size is given as the distance to the pose which is calculated and is either one or three additional pose(s) before and after e.g., spanning three poses in total or spanning seven poses in

<sup>4</sup><https://github.com/anonPETs2023/PETs2023>

total, respectively. This strategy follows a similar idea to low-pass filtering, as it retains the main movement but removes detailed deviations.

*Coarsening macro:* As we were interested in the most significant information of the walker position, we removed the least significant part of each marker position in a pose for all poses. The effect is that the grid on which the walker moves is becoming more coarse. We removed all digits either below the thousandth (1000) or the hundredth digit (100).

*Remove body parts:* We measured how much an individual body part (head, torso, hip, arms, legs) contributes to the overall recognition performance. This was done by removing the body part from the data by setting its marker positions to zero.

**3.3.2 Micro Features.** The micro features are the counterparts to the macro features. Here we kept the small variations of the gait cycle and the least significant parts of the marker positions.

*Remove trajectories:* Contrasting remove variations, we removed the ideal marker trajectories from the data by calculating the ideal trajectory as described in remove variation via either rolling average or interpolation with a window size of 1 or 3. The ideal trajectory was then subtracted from the real trajectory, which leaves us with the distances of the ideal marker positions to the real ones. This strategy resembles high-pass filtering, as it removes the main movement and only retains the minor specifics of the current sample.

*Coarsening micro:* We eliminated the most significant part of the walker positions by removing the most significant parts of each marker position value. We removed all digits above the hundredth (100), tenth (10), or first digit (1) position

*Keep body part:* We measured how much recognition performance the individual body parts have alone without the rest of the body. All remaining other body parts are set to zero.

*Amplitude/Frequency equalization:* The walking amplitude and frequency were equalized between all individuals to perturb their influence on the recognition. Informed by previous studies [49], we calculated a gait representation of each individual by using the average pose, the first four components of a principal component analysis (PCA), and a sinus function fit on these components to represent the gait cycle of a person. We then equalized the frequency or amplitude of the fitted sinus function by means of the group-level average.

**3.3.3 Static Features.** The static features capture the time-invariant features of the walker by removing the dynamic part of the gait motion. We therefore kept the proportions of the walker.

*Static pose:* We used only an average pose or the first pose of each sample, thus removing the dynamic component of the gait data.

*Resampling:* We downsampled the data to 10 frames, and therefore removed most of the dynamic content from the data.

**3.3.4 Dynamic Features.** The dynamic features are the counterpart to the static features and aim to only retain the dynamic part of the motion.

*Motion extraction:* Instead of using the individual poses, we used their difference (i.e., keeping only the variations between poses) and hence removed the static features.

*Normalization:* We normalized the static features in a sequence by either normalizing the height axis (y-axis), all axes or normalizing each dimension over the entire sequence of poses.

**3.3.5 Combinations of Perturbations.** Besides evaluating each of the features alone, we also investigated their combinations. Two perturbation techniques were combined by applying them sequentially to the data. Due to some techniques (first pose, average pose) not returning a time series, not all combinations of methods are possible. As the overall number of combinations is quite high, we focused on representatives of each class of features. We picked those representatives by their anonymization impact on the data.

## 3.4 Recognition System

To test the impact of omitting features from human gait, and hence their contribution to inference, we implemented a gait recognition system. It is based on the system by Horst et al. [23] using Python 3.8.3 [43], Scikit-learn 0.23.1 [41], and NumPy 1.18.5 [19]. We used two feature vectors to represent a data sample: *flatten* which concatenates all poses of a sample into a single vector, and *reduced angles* which first calculates a reduced representation of 17 markers representing the main body parts and then calculated 10 joint angles from this representation.

Next, the data was split into train (75%) and test (25%) data. Here we differentiated between the identity and sex recognition. For identity recognition, we split the samples for each identity so that we have every identity in both sets. While for sex recognition we split the samples identity-wise, making sure that every identity is only in one of the sets. We did so to make sure that the classifier cannot learn the identity to perform sex recognition. Following the split, we then scaled the data in each set by subtracting the mean and then scaling with the standard deviation before we performed a principal component analysis (PCA) to reduce the dimensions of the samples. As a classifier, we used a support vector machine (SVM) using a radial basis function (RBF) kernel. For the training of the SVM we used 10-fold cross-validation with the train set before we tested the best performing model on the test set. In order to account for the random splitting of the data, we ran the entire process 10 times.

## 3.5 Utility

Besides investigating the identity and sex recognition performance of our features we also sought to understand how much the features contribute to the utility of envisioned applications. As our use case (see Sec. 2.1) is to transfer the gait motion onto a digital avatar, the goal is to retain as much naturalness in the motion data as possible. In order to measure the corresponding effect, we performed an online survey with 22 human participants (13 male, age: 18–60 years) which we asked to rate the naturalness of the perturbed gait sequences. Participants were shown renderings of two gait sequences for each perturbation in which the walkers (one male and one female walker, individually) were shown from the side 45 degrees rotated around the z-axis towards the camera. The renderings are identical to the example videos we provided in Section 3.3. All sequences were shown in random order. The participants then rated on a scale from 1 (worst) to 5 (best) how natural looking the gait sequence appeared to them. The survey data collection is under

the umbrella of the project ("Privatsphäre von Körperbewegungen") approved on 30.09.2021 by the ethical committee of KIT and was conducted in accordance with the Declaration of Helsinki. The survey data was collected in anonymized form.

## 4 RESULTS

In this section we present the results of our obfuscation experiments, by reporting the recognition performance of the chosen feature categories. The results for identity and sex recognition in two contrasting feature categories (macro vs. micro, dynamic vs. static) are reported each. Note, that we report the body part removals (body parts vs. rest body) separate from the macro and micro features for easier comparability.

As we conducted recognition experiments and the classified classes (for both identity and sex recognition) have nearly the same number of samples per class, we selected accuracy as our metric. Accuracy is defined as the number of correctly classified samples divided by the number of all classified samples. In Section 3 we described the two feature vectors we used in our recognition system. Since we were interested in how much identifiable information remains in the data after the perturbation has been applied, we always report the best performing feature vector.

### 4.1 Macro vs. Micro Features

We start by comparing macro to micro features. For both, identity (Fig. 5) and sex recognition (Fig. 6), we can see similar effects for the macro features: The variation removal via rolling average and interpolation shows no effect on the accuracy. The coarsening of all digits below the 100th digit has no effect, while coarsening from the 1000th digit position leads to a drop in accuracy for sex recognition to 91% and identity recognition to 77%. For the micro features, we see a difference between identity and sex recognition. Only trajectory removal using an interpolation window of 1 drops the accuracy of the identity recognition, while all of the others lead to a drop in sex recognition to about 90%. For the micro coarsening methods we again see that identity recognition is not affected by coarsening everything higher than the 100th digit, while for sex recognition we see a drop of accuracy to 90%. Then coarsening the digits above the 10th digit leads for both, identity and sex recognition, to chance level accuracy. The results show that sex recognition is more dependent on the macro feature than on the micro features, while the identity can be perfectly inferred from both of them.

### 4.2 Individual Body Parts in Isolation vs. Reduced Whole Bodies

Next, we evaluate perturbations of individual isolated body parts in contrast to reduced whole body configurations (i.e., certain body parts were removed) (see Fig. 7 and Fig. 8). On the one hand, only the specified body part is used for the recognition ("keep"), while on the other hand, the whole body minus the specific body part ("remove") is employed. Fig. 7 shows that the removal of the legs slightly reduces the identity recognition accuracy to 97%. At the same time, it is the only body part that achieves 100% recognition accuracy alone. In contrast, keeping only the head as the standalone body part achieves the strongest prevention from identity recognition,

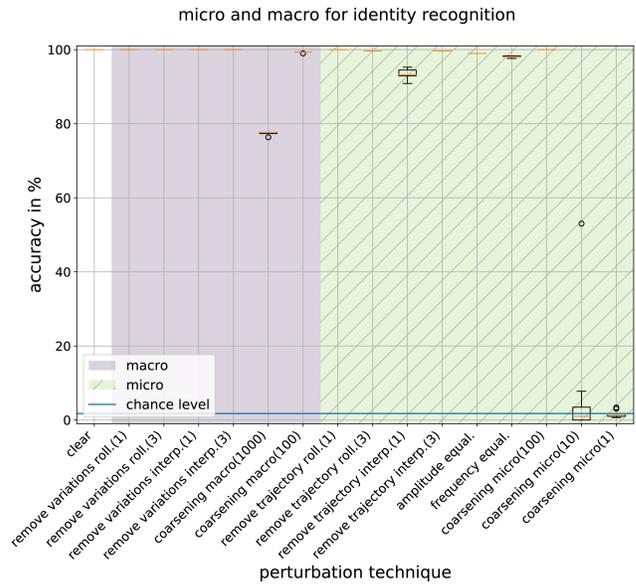


Figure 5: Boxplot of accuracy results for micro and macro features for identity recognition given in percent.

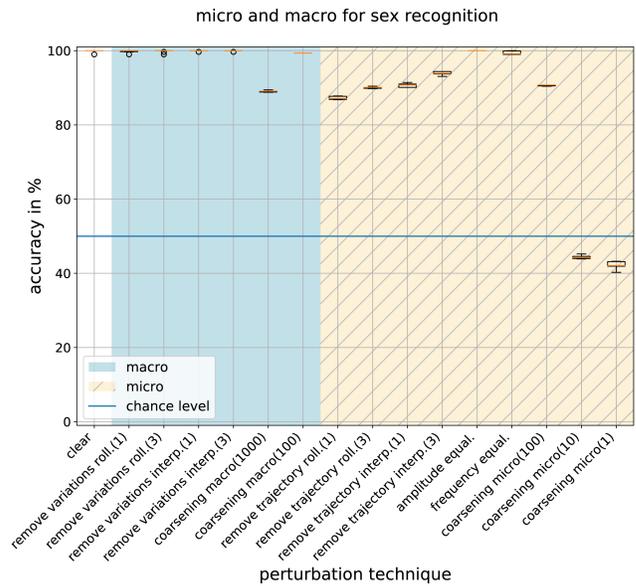
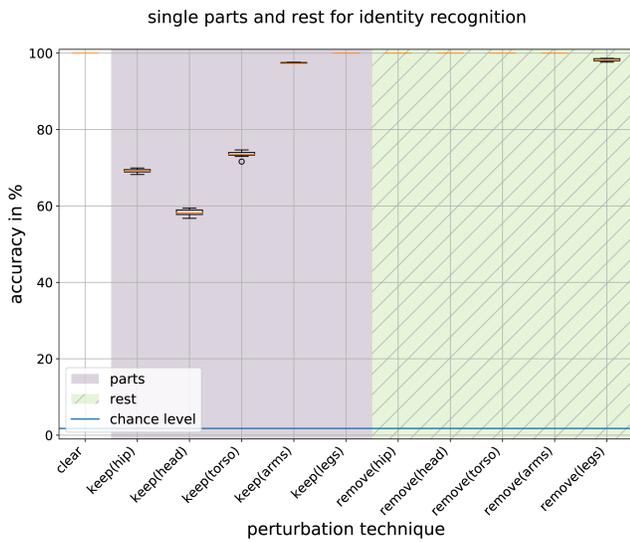


Figure 6: Boxplot of accuracy results for micro and macro features for sex recognition given in percent.

reducing the accuracy to less than 60%. Only slightly improved performance is achieved by the standalone body parts torso or hip.

Results for sex recognition with respect to perturbed body parts are displayed in Fig. 8. We find the same small reduction in accuracy for the removal of the legs as we saw for identity recognition, while it is again the only body part to achieve the full recognition accuracy as a standalone body part. However, for the other body parts, we find that their removal does not impact the sex recognition



**Figure 7: Accuracy boxplots of results for individual body parts and all of the remaining body parts for identity recognition given in percent.**

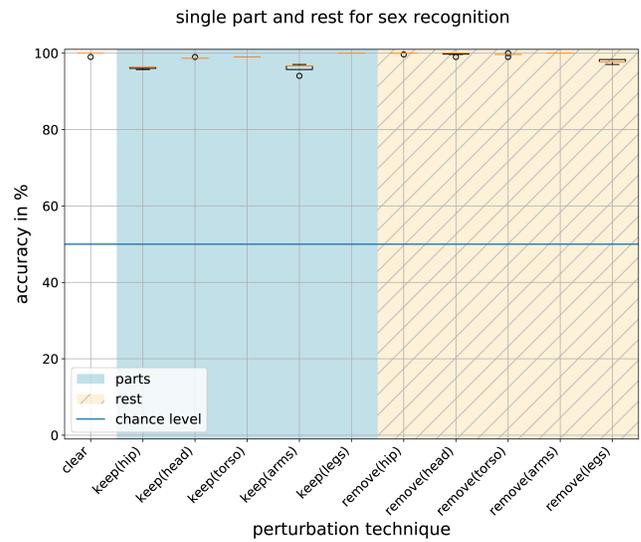
score. Additionally, our data shows only small effects on using only individual body parts in isolation. Comparing identity to sex recognition, head, hip, and torso alone fare much better for sex than for identity recognition. These results suggest that even the limited form information which is integrated over time into dynamic form information is sufficient to identify biological traits such as sex or even identity. This finding is in line with human perception research. For example, Kozłowski et al. [28] found that longer strides are perceived as more masculine. Center of moment contains sex information (see also Section 2.3; [28, 37]). That is, as long as the stimuli contains information about certain body parts, sex and identity recognition is possible.

### 4.3 Dynamic vs. Static Features

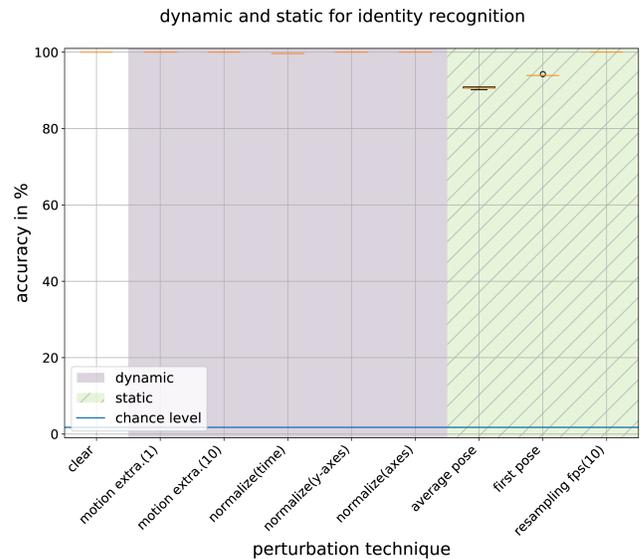
Thirdly, we investigate the effects of dynamic and static feature perturbation on recognition performance. In the case of identity recognition, depicted in Fig. 9, we observe that only using the average pose or the first pose reduces the recognition accuracy slightly to 91% and 94% respectively, while other feature manipulations show no effect on identity recognition. For sex recognition (Fig. 10), our results show that while static features have close to no effect on accuracy, all dynamic features appear to do so. So we can conclude that the static features are more important for the sex recognition than the dynamic ones.

### 4.4 Combination of Features

Here, we evaluate the combination of selected perturbation techniques from each category. Due to the further removal of data, we expected to see larger reductions in the classification accuracy for both identity and sex recognition. We also expected that with fewer data available the classification process becomes more unsteady and



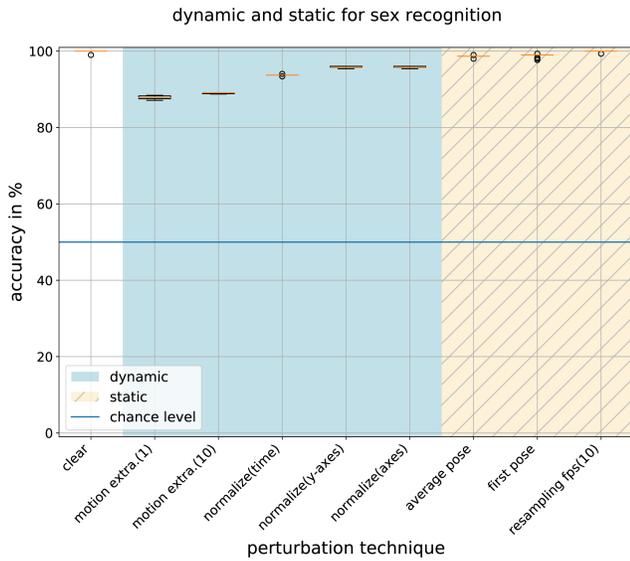
**Figure 8: Accuracy boxplots of results for individual body parts and all of the remaining body parts for sex recognition given in percent.**



**Figure 9: Boxplots of accuracy results for dynamic and static features for identity recognition given in percent.**

therefore the variance between the results will be larger. Further, the reduction of data can lead to a simplification of the data, which then is easier to classify.

The combination of body parts head and legs with the static, dynamic, micro, and macro categories for sex recognition are shown in Fig. 8. Most of the legs combinations remain at 100% accuracy. Only in combination with average pose and coarsening micro (100), a slight decrease in accuracy can be observed. When the legs are combined with coarsening macro (1000) we observe a large decrease in

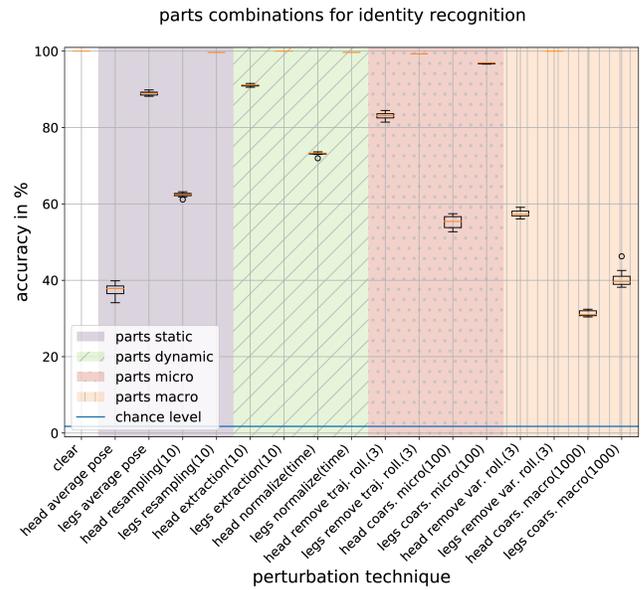


**Figure 10: Accuracy boxplots of results for dynamic and static features for sex recognition given in percent.**

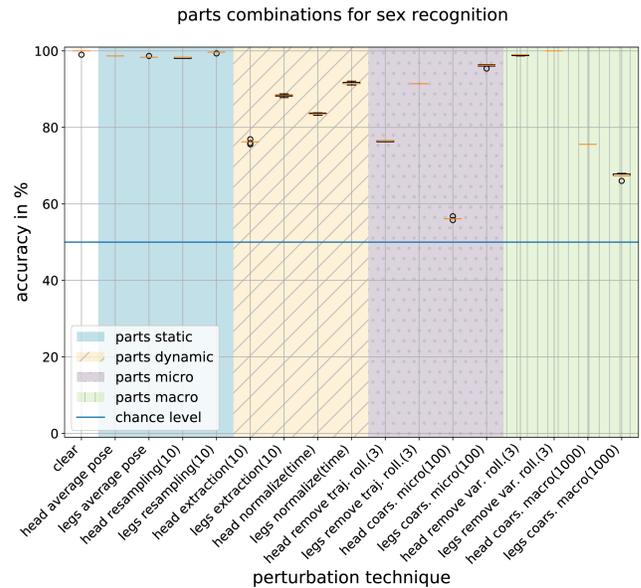
accuracy to close to 40%, while both of these perturbations alone do not have an effect on the accuracy. The head (head alone achieves 60% identity recognition) combinations are more of a mixed bag. While average pose and coarsening macro further reduce the accuracy; resampling, coarsening micro, and remove variations do not have an additional effect on the accuracy. However, motion extraction, time normalization, and remove trajectory lead to an increase in the recognition accuracy. All three methods focus more on the smaller variations in the data, providing an indication that the identification of individuals via their head motion is more dependent on the dynamic parts than the general movement.

Focusing on the combinations with head and legs for sex recognition, we find that while there is no effect on accuracy in combination with static features, the combination with dynamic features has deleterious effects on accuracy. Specifically, the combination of head and motion extraction results in a drop of accuracy to 75%. For the micro combinations, we again find that the combinations with the head suffer the largest accuracy reduction. Here the combination with micro coarsening nearly reaches chance level, while the same combination with the legs stays above 90%. When we compare this with the macro coarsening of the macro features, we find that the legs drop to a lower accuracy than the head. This leads us to conclude that the sex recognition via the head data is much more dependent on the macro part of the head position, while the sex recognition via the legs depends more on the micro part of the positions.

The results of macro, micro, dynamic, and static feature combinations for identity recognition are shown in Fig. 13. The macro-static combinations show accuracy decreases for the combinations that contain macro coarsening. We also see these decreases when we look at the combination of macro dynamic features in which the macro coarsening leads to a decrease in performance. The last combinations show an accuracy decrease in the removal of the



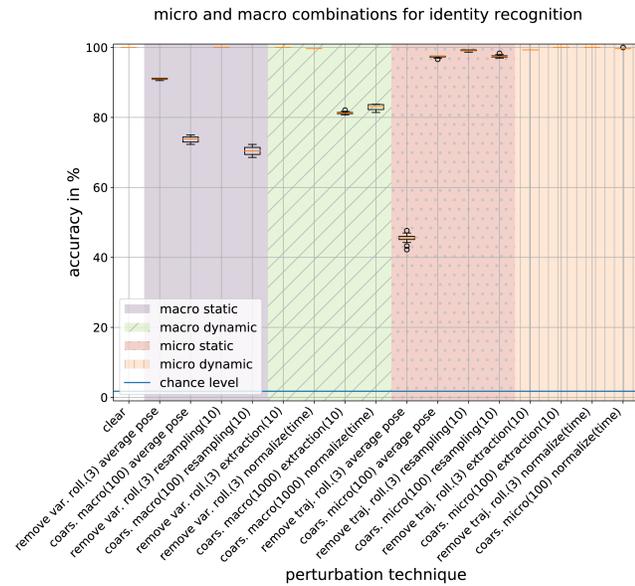
**Figure 11: Accuracy boxplots of results for legs and head in combination with the other categories for identity recognition given in percent.**



**Figure 12: Accuracy boxplots of results for legs and head in combination with the other categories for sex recognition given in percent.**

trajectory plus average pose which drops the recognition accuracy to 45%. Comparing the removal of variations and trajectory in combination with the average pose, show that the general trajectory of the walker contains much identifiable information in their overall

characteristic, while the small variations from the trajectories are only meaningful when their dynamic features are preserved.

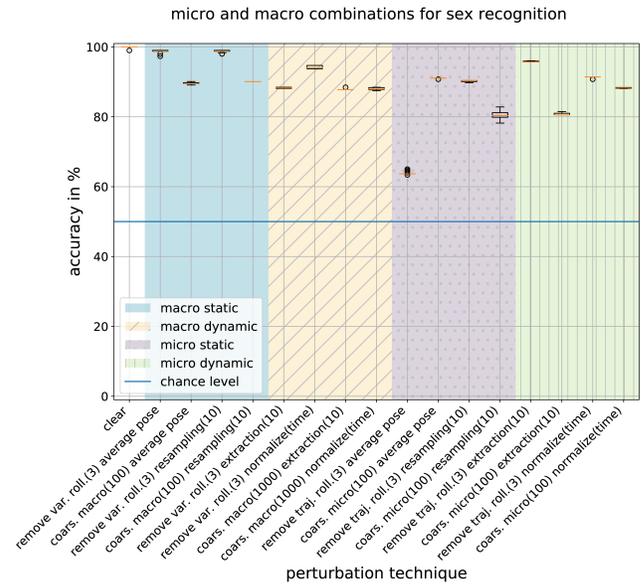


**Figure 13: Accuracy boxplots of results for macro, micro, dynamic, and static combinations for identity recognition given in percent.**

Lastly, we look at the same feature combinations as before but this time for sex recognition (see Fig. 14). In the case of the combination of macro and static features, the removal of the variations does not lead to a drop in accuracy, while both combinations with coarsening macro drop to the same accuracy level of about 90%. This suggests that obfuscations in combination with macro features have a bigger impact on the accuracy in comparison to combinations with static features. All macro-dynamic combinations result in a decrease of performance to about 90%. Furthermore, removing the variations plus macro coarsening increases the performance slightly when compared to just performing the same macro coarsening alone (see Fig. 6). In the micro-static combinations, we find the removal of the trajectory average pose combination to create a large drop in accuracy.

### 4.5 Utility

Finally, we report the results of the naturalness evaluation for all perturbation techniques that have a median rating score which is greater than 1 (all techniques that retain some utility; on the 1-5 scale described in Sect. 3.5) and are shown in Fig. 15. Perturbations that resulted in a median score below that, were assumed to retain no utility and are therefore not plotted. First, we note that none of the micro feature perturbations retained any gait naturalness. In the static category only average and first pose managed to appear minimally natural, with median naturalness scores of 2. The exclusion of body parts of the walkers had deleterious effects on the perceived naturalness, while still maintaining some level of naturalness depending on the specific removed body part. Interestingly, keeping

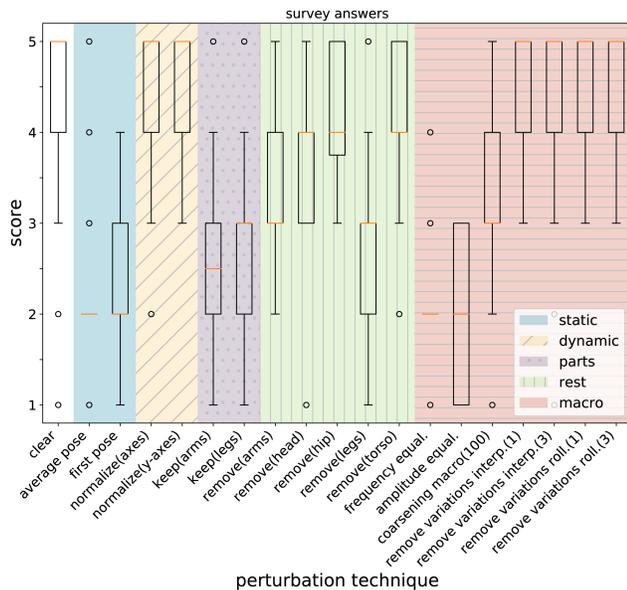


**Figure 14: Accuracy boxplots of results for macro, micro, dynamic, and static combinations for sex recognition given in percent.**

only the arms or legs of the walker was rated as still somewhat natural, whereas all other individual body parts in isolation were rated as non-natural. The normalization of all axes and the normalization of the y-axes achieve the same level (median of 5) of naturalness as the clear data. The only other techniques that achieve the same naturalness ratings are the remove variation techniques. In general, these results are within our expectations, as perturbing the data should either maintain the same level of naturalness or decrease it. The fact that most of the macro features retained the naturalness of the walker is also unsurprising, as they preserve the majority of the gait variations while the small variations we kept in the micro features are not perceived as natural anymore. We did not evaluate the naturalness of the combinations, however, we assume that a combination will at most reach the minimum naturalness rating score of its two used perturbation techniques.

## 5 DISCUSSION

Using ML for gait recognition based on motion capture data, we investigated the importance of features based on findings in psychology for identity and sex recognition. The findings reported here, suggest that all of the features reported by psychology are transferable to ML approaches in identification performance based on walking motion. The identification procedure is robust as even when large parts of the data are removed the identification rates are high, only when multiple features are removed from the data a significant impact on the accuracy can be observed. Consistent with previous studies in psychology and neuroscience [42, 53, 49], we found that dynamic and static features contain much identifiable information, hinting at strong temporal and physiological dependencies in the data. We anticipate that for the development



**Figure 15: Boxplots of the naturalness rating scores for the perturbation techniques that retained utility.**

of suitable anonymization techniques for gait data the dependencies between the features have to be accounted for, as otherwise, the reconstruction of the clear data is likely possible. For example, noise that is applied to the marker positions could be removed by smoothing the trajectories, or missing markers can be reconstructed from the position of the remaining ones. Wang et al. [52] have convincingly demonstrated this, showing how adding noise does not effectively perturb correlated data. Interestingly, the removal of body parts and the subsequent performance accuracy alone indicates a high redundancy in the data, and as such focusing on a single feature for anonymization is unlikely to achieve a meaningful anonymization effect. This effect, albeit in a much weaker form, has previously been shown in human person and biological motion perception studies: The elimination of some local information, for example by removing PLD dots corresponding to body parts, does not affect the recognition as long as a certain degree of global form revealing dynamic posture changes is preserved [6, 29].

Both, the overall trajectories of the gait as well as small variations in the data, allow for recognition of individuals. Thus, making it necessary to adjust the overall gait trajectory for anonymization purposes. The overall pattern of results here provides converging evidence for the need to consider gait motion capture a strong personal identifiable trait, even when recorded at low resolution or low frame rate. Many features, as investigated here — macro, micro, dynamic and static features as well as individual body parts — contain strong identifiable information about both, the identity and the sex of a human walker. With our simple ML-based feature perturbation approach we found that coarsening the marker positions precision, with the respective recognition performances of 45% and 2% for sex and identity exhibited the strongest reduction of classification accuracy while removing dynamic & static features generally only reduced recognition slightly. However, our utility

evaluation of the features shows that the perceived naturalness of the perturbed data is diminished when the general motion or body structure of the walkers is removed. Thus we see a strong indication that in order to develop strong anonymization for gait data, while keeping its utility intact, a holistic approach is required. Such an approach should take the dependencies in the data and the requirement for natural-looking results into account, for example by generating synthetic gait trajectories.

## 5.1 Limitations and Future Work

The present work is based on a data set of 57 young adult individuals and as such it might be possible to achieve superior anonymization results for larger sample sizes. However, as we have shown gait data does contain a large amount of identifiable information, so larger effects from bigger samples are unlikely. The present work presents results on one sole gait cycle per sample, future work should include multiple sequential gait cycles or gait data from multiple sessions. Furthermore, all individuals were from a similar age cohort, including different lifespan age brackets or longitudinal data might lead to more meaningful and representative results. However, we believe that having a cohort of very similar individuals also strengthens the recognition results, as it becomes more difficult to tell the individuals apart. It is possible that with the improvements of machine learning approaches, better classification results can be achieved on our perturbed data. As such our approach only gives a lower bound how much identifiable information remains in the perturbed data. This fact is also shown by some of the combinations of perturbation techniques where the combinations achieved higher recognition accuracy than the individual techniques alone.

With regards to the user study we would like to point out that our definition of utility only takes into account how natural other people perceived the anonymized PLD gait sequences shown to them. We did not investigate if the original walker themselves would find their perturbed gait to be natural. We did so because we assumed that the device used in our system-and-threats-model is trusted by the user and therefore would display the real gait (pre-transfer to the service provider described in Sect. 2.1; labeled "clear" in the present work) to the user as it is recorded locally in real-time, instead of an anonymized version of the user's gait. Furthermore, we based our present investigation on an existing open-source dataset and therefore have no access in an ethical and legal way to the original walkers due to inter alia data protection and privacy reasons. Future studies that obtain their own motion capture recordings could include an evaluation of utility by asking the recorded walkers themselves to evaluate their perturbed gait or other movements recorded with motion capture.

For additional future work we propose to conduct the same set of experiments with human observers to directly compare human and machine gait recognition, in order to gain insight into how both differ in regards to identifying individuals and their sex. Although, the human ability to process biological motion such as gait-based person perception and recognition is susceptible to viewer-specific influences such as age [8], social factors (e.g., interpersonal context, stereotypes) [9, 26], neurodevelopmental disorders (e.g., autism, schizophrenia) [9], and other potential experimental, concomitant, and individual factors [17, 20, 50, 11, 26, 30]. Thus, utilizing machine

gait recognition provides a more objective evaluation method for different anonymization techniques.

## 6 CONCLUSION

In this paper, we address the question of how much specific features of human gait contribute to the ability to discern the identity or sex of different human individuals in gait data. Here, we found that overall identification performance was indeed very robust. Removing large parts of the data, either by omitting body parts or reducing spatial and temporal resolution, did have little effect on the recognition performance.

One possible interpretation of the findings is that gait is idiosyncratic and very redundant. Moreover, gait can be considered an individual trait that shows little variability over time and even lifespan. Studies reported that major adult gait emerges already at the age of five years, although age-related effects such as slower gait or shorter steps as well as age-related body proportion changes have been found as well [37].

Our results suggest that gait will be very hard to anonymize effectively. This entails that anonymization cannot be achieved with simple means, but will require intricate approaches that take the inter-dependency of the connected body, as well as the overall generating process of the walking human into consideration. Utility can only be retained when the macro structure of the walker and its dynamic are largely kept intact.

## ACKNOWLEDGMENTS

This work has been supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden, and by funding of the Helmholtz Association through the KASTEL Security Research Labs (HGF Topic 46.23).

## REFERENCES

- [1] David Alvarez Melis and Tommi Jaakkola. 2018. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors. Volume 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/3e9f0fc9b2f89e043bc6233994dfcf76-Paper.pdf>.
- [2] Vijay Arya, Rachel K. E. Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Q. Vera Liao, Ronny Luss, Aleksandra Mojsilović, Sami Mourad, Pablo Pedemonte, Ramya Raghavendra, John Richards, Prasanna Sattigeri, Karthikeyan Shanmugam, Moninder Singh, Kush R. Varshney, Dennis Wei, and Yunfeng Zhang. 2019. One explanation does not fit all: a toolkit and taxonomy of AI explainability techniques. *arXiv*. DOI: <https://doi.org/10.48550/arXiv.1909.03012>.
- [3] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10, 7, e0130140. ISSN: 1932-6203. DOI: <https://doi.org/10.1371/journal.pone.0130140>.
- [4] Michal Balazia and Petr Sojka. 2018. Gait recognition from motion capture data. *ACM Trans. Multimedia Comput. Commun. Appl.*, 14, 1, 1–18. ISSN: 1551-6857, 1551-6865. DOI: <https://doi.org/10.1145/3152124>.
- [5] Michal Balazia and Petr Sojka. 2017. You are how you walk: uncooperative mocap gait identification for video surveillance with incomplete and noisy data. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, 208–215. DOI: <https://doi.org/10.1109/BTAS.2017.8272700>.
- [6] JA Beintema and Markus Lappe. 2002. Perception of biological motion without local image motion. *Proceedings of the National Academy of Sciences*, 99, 8, 5661–5663. DOI: <https://doi.org/10.1073/pnas.082483699>.
- [7] Emmanuelle Bellot, Etienne Abassi, and Liuba Papeo. 2021. Moving toward versus away from another: how body motion direction changes the representation of bodies and actions in the visual cortex. *Cerebral Cortex*. ISSN: 1047-3211. DOI: <https://doi.org/10.1093/cercor/bhaa382>.
- [8] Jutta Billino and Karin S Pilz. 2019. Motion perception as a model for perceptual aging. *Journal of vision*, 19, 4, 3–3. DOI: <https://doi.org/10.1167/19.4.3>.
- [9] Danielle Z Bolling, Kevin A Pelphrey, and Martha D Kaiser. 2013. Social inclusion enhances biological motion processing: a functional near-infrared spectroscopy study. *Brain topography*, 26, 2, 315–325. DOI: <https://doi.org/10.1007/s10548-012-0253-y>.
- [10] Talia Brandman and Galit Yovel. 2016. Bodies are represented as wholes rather than their sum of parts in the occipital-temporal cortex. *Cerebral Cortex*, 26, 2, 530–543. ISSN: 1460-2199. DOI: [10.1093/cercor/bhu205](https://doi.org/10.1093/cercor/bhu205).
- [11] Patrick Connor and Arun Ross. 2018. Biometric recognition by gait: a survey of modalities and features. *Computer vision and image understanding*, 167, 1–27. DOI: <https://doi.org/10.1016/j.cviu.2018.01.007>.
- [12] James E Cutting and Lynn T Kozlowski. 1977. Recognizing friends by their walk: gait perception without familiarity cues. *Bulletin of the psychonomic society*, 9, 5, 353–356. DOI: <https://doi.org/10.3758/BF03337021>.
- [13] Silvia Del Din, Alan Godfrey, and Lynn Rochester. 2016. Validation of an accelerometer to quantify a comprehensive battery of gait characteristics in healthy older adults and parkinson's disease: toward clinical and at home use. *IEEE Journal of Biomedical and Health Informatics*, 20, 3, 838–847. DOI: <https://doi.org/10.1109/JBHI.2015.2419317>.
- [14] John David N. Dionisio, William G. Burns III, and Richard Gilbert. 2013. 3d virtual worlds and the metaverse: current status and future possibilities. *ACM Comput. Surv.*, 45, 3, (July 2013). ISSN: 0360-0300. DOI: <https://doi.org/10.1145/2480741.2480751>.
- [15] Bjoern M. Eskofier, Peter Federolf, Patrick F. Kugler, and Benno M. Nigg. 2013. Marker-based classification of young-elderly gait pattern differences via direct PCA feature extraction and SVMs. *Computer Methods in Biomechanics and Biomedical Engineering*, 16, 4, 435–442. ISSN: 1025-5842, 1476-8259. DOI: <https://doi.org/10.1080/10255842.2011.624515>.

- [16] Ruth C. Fong and Andrea Vedaldi. 2017. Interpretable explanations of black boxes by meaningful perturbation. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 2017 IEEE International Conference on Computer Vision (ICCV). IEEE, Venice, (October 2017), 3449–3457. ISBN: 978-1-5386-1032-9. DOI: <https://doi.org/10.1109/ICCV.2017.371>.
- [17] Celia Foster, Mintao Zhao, Javier Romero, Michael J Black, Betty J Mohler, Andreas Bartels, and Isabelle Bühlhoff. 2019. Decoding subcategories of human bodies from both body- and face-responsive cortical regions. *NeuroImage*, 202, 116085. DOI: <https://doi.org/10.1016/j.neuroimage.2019.116085>.
- [18] Marco Gandolfo and Paul E Downing. 2020. Asymmetric visual representation of sex from human body shape. *Cognition*, 205, 104436. ISSN: 0010-0277. DOI: <https://doi.org/10.1016/j.cognition.2020.104436>.
- [19] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. Array programming with NumPy. *Nature*, 585, 7825, 357–362. DOI: <https://doi.org/10.1038/s41586-020-2649-2>.
- [20] John D Herrington, Charlotte Nymberg, and Robert T Schultz. 2011. Biological motion task performance predicts superior temporal sulcus activity. *Brain and cognition*, 77, 3, 372–381. DOI: <https://doi.org/10.1016/j.bandc.2011.09.001>.
- [21] Yuki Hirose, Kazuaki Nakamura, Naoko Nitta, and Noboru Babaguchi. 2019. Anonymization of gait silhouette video by perturbing its phase and shape components. In *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. 2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC). IEEE, Lanzhou, China, (November 2019), 1679–1685. ISBN: 978-1-72813-248-8. DOI: <https://doi.org/10.1109/APSIPAASC47483.2019.9023196>.
- [22] Fabian Horst, Sebastian Lapuschkin, Wojciech Samek, Klaus-Robert Müller, and Wolfgang I Schöllhorn. 2018. A public dataset of overground walking kinetics and full-body kinematics in healthy individuals. *Mendeley Data Repository*, 1. DOI: <http://dx.doi.org/10.17632/svx74xcrjr.1>.
- [23] Fabian Horst, Sebastian Lapuschkin, Wojciech Samek, Klaus-Robert Müller, and Wolfgang I Schöllhorn. 2019. Explaining the unique nature of individual gait patterns with deep learning. *Scientific reports*, 9, 1, 1–13. DOI: <https://doi.org/10.1038/s41598-019-38748-8>.
- [24] Marina Ivacic-Kos, Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas. 2014. Person de-identification in activity videos. In *2014 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. 2014 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO). IEEE, Opatija, Croatia, (May 2014), 1294–1299. ISBN: 978-953-233-077-9 978-953-233-081-6. DOI: <https://doi.org/10.1109/MIPRO.2014.6859767>.
- [25] Gunnar Johansson. 1973. Visual perception of biological motion and a model for its analysis. *Perception & psychophysics*, 14, 2, 201–211. ISSN: 0031-5117. DOI: <https://doi.org/10.3758/BF03212378>.
- [26] Kerri L Johnson, Lawrie S McKay, and Frank E Pollick. 2011. He throws like a girl (but only when he’s sad): emotion affects sex-decoding of biological motion displays. *Cognition*, 119, 2, 265–280. ISSN: 0010-0277. DOI: <https://doi.org/10.1016/j.cognition.2011.01.016>.
- [27] Łukasz Kidziński, Bryan Yang, Jennifer L Hicks, Apoorva Rajagopal, Scott L Delp, and Michael H Schwartz. 2020. Deep neural networks enable quantitative movement analysis using single-camera videos. *Nature communications*, 11, 1, 1–10. DOI: <https://doi.org/10.1038/s41467-020-17807-z>.
- [28] Lynn T Kozlowski and James E Cutting. 1977. Recognizing the sex of a walker from a dynamic point-light display. *Perception & psychophysics*, 21, 6, 575–580. ISSN: 1532-5962. DOI: <https://doi.org/10.3758/BF03198740>.
- [29] Joachim Lange, Karsten Georg, and Markus Lappe. 2006. Visual perception of biological motion by form: a template-matching analysis. *Journal of vision*, 6, 8, 6–6. DOI: <https://doi.org/10.1167/6.8.6>.
- [30] Florian Loffing, Florian Sölter, Norbert Hagemann, and Bernd Strauss. 2016. On-court position and handedness in visual anticipation of stroke direction in tennis. *Psychology of Sport and Exercise*, 27, 195–204. DOI: <https://doi.org/10.1016/j.psychsport.2016.08.014>.
- [31] Elena Lorenzi, Uwe Mayer, Orsola Rosa-Salva, and Giorgio Vallortigara. 2017. Dynamic features of animate motion activate septal and preoptic areas in visually naïve chicks (*gallus gallus*). *Neuroscience*, 354, 54–68. DOI: <https://doi.org/10.1016/j.neuroscience.2017.04.022>.
- [32] Fani Loula, Sapna Prasad, Kent Harber, and Maggie Shiffrar. 2005. Recognizing people from their movement. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1, 210. ISSN: 1939-1277. DOI: <https://doi.org/10.1037/0096-1523.31.1.210>.
- [33] Valeria Manera, Ben Schouten, Cristina Becchio, Bruno G Bara, and Karl Verfaillie. 2010. Inferring intentions from biological motion: a stimulus set of point-light communicative interactions. *Behavior research methods*, 42, 1, 168–178. ISSN: 1554-351X. DOI: <https://doi.org/10.3758/BRM.42.1.168>.
- [34] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. 2018. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21, 9, 1281–1289. DOI: <https://doi.org/10.1038/s41593-018-0209-y>.
- [35] Richard Matovu, Abdul Serwadda, David Irakiza, and Isaac Griswold-Steiner. 2018. Jekyll and hyde: on the double-faced nature of smart-phone sensor noise injection. In *2018 International Conference of the Biometrics Special Interest Group (BIOSIG)*. 2018 International Conference of the Biometrics Special Interest Group (BIOSIG). IEEE, Darmstadt, (September 2018), 1–6. ISBN: 978-3-88579-676-3. DOI: <https://doi.org/10.23919/BIOSIG.2018.8553043>.

- [36] Sina Mehdizadeh, Hoda Nabavi, Andrea Sabo, Twinkle Arora, Andrea Iaboni, and Babak Taati. 2021. Concurrent validity of human pose tracking in video for measuring gait parameters in older adults: a preliminary analysis with multiple trackers, viewing angles, and walking directions. *Journal of NeuroEngineering and Rehabilitation*, 18, 1, 1–16. DOI: <https://doi.org/10.1186/s12984-021-00933-0>.
- [37] Joann M Montepare and Leslie Zebrowitz-McArthur. 1988. Impressions of people created by age-related qualities of their gaits. *Journal of personality and social psychology*, 55, 4, 547. ISSN: 1939-1315. DOI: <https://doi.org/10.1037/0022-3514.55.4.547>.
- [38] Carlos Morato, Krishnanand N Kaipa, Boxuan Zhao, and Satyandra K Gupta. 2014. Toward safe human robot collaboration by using multiple kinects based real-time human tracking. *Journal of Computing and Information Science in Engineering*, 14, 1. DOI: <https://doi.org/10.1115/1.4025810>.
- [39] Edward R Morrison, Hannah Bain, Louise Pattison, and Hannah Whyte-Smith. 2018. Something in the way she moves: biological motion, body shape, and attractiveness in women. *Visual Cognition*, 26, 6, 405–411. DOI: <https://doi.org/10.1080/13506285.2018.1471560>.
- [40] Liuba Papeo, Moritz F Wurm, Nikolaas N Oosterhof, and Alfonso Caramazza. 2017. The neural representation of human versus nonhuman bipeds and quadrupeds. *Scientific reports*, 7, 1, 1–8. ISSN: 2045-2322. DOI: <https://doi.org/10.1038/s41598-017-14424-7>.
- [41] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. 2011. Scikit-learn: machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830. <http://scikit-learn.sourceforge.net>.
- [42] Yujia Peng, Hannah Lee, Tianmin Shu, and Hongjing Lu. 2020. Exploring biological motion perception in two-stream convolutional neural networks. *Vision Research*, 178, 28–40. ISSN: 0042-6989. DOI: <https://doi.org/10.1016/j.visres.2020.09.005>.
- [43] Python Core Team. 2019. *Python: A dynamic, open source programming language*. Python version 3.8.3. Python Software Foundation. <https://www.python.org/>.
- [44] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. 2019. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller, editors. *Lecture Notes in Computer Science*. Volume 11700. Springer International Publishing, Cham. ISBN: 978-3-030-28953-9 978-3-030-28954-6. DOI: <https://doi.org/10.1007/978-3-030-28954-6>.
- [45] Shiqi Yu, Tieniu Tan, Kaiqi Huang, Kui Jia, and Xinyu Wu. 2009. A study on gait-based gender classification. *IEEE Transactions on Image Processing*, 18, 8, 1905–1910. ISSN: 1057-7149, 1941-0042. DOI: <https://doi.org/10.1109/TIP.2009.2020535>.
- [46] Noa Simhi and Galit Yovel. 2020. Dissociating gait from static appearance: a virtual reality study of the role of dynamic identity signatures in person recognition. *Cognition*, 205, 104445. ISSN: 0010-0277. DOI: <https://doi.org/10.1016/j.cognition.2020.104445>.
- [47] Djordje Slijepcevic, Fabian Horst, Sebastian Lapuschkin, Anna-Maria Raberger, Matthias Zeppelzauer, Wojciech Samek, Christian Breiteneder, Wolfgang I. Schöllhorn, and Brian Horschak. 2020. On the explanation of machine learning predictions in clinical gait analysis. *arXiv*. DOI: <https://doi.org/10.48550/arXiv.1912.07737>.
- [48] Ngoc-Dung T. Tieu, Huy H. Nguyen, Hoang-Quoc Nguyen-Son, Junichi Yamagishi, and Isao Echizen. 2017. An approach for gait anonymization using deep learning. In *2017 IEEE Workshop on Information Forensics and Security (WIFS)*. 2017 IEEE Workshop on Information Forensics and Security (WIFS). IEEE, Rennes, France, (December 2017), 1–6. ISBN: 978-1-5090-6769-5. DOI: <https://doi.org/10.1109/WIFS.2017.8267657>.
- [49] Nikolaus F Troje, Cord Westhoff, and Mikhail Lavrov. 2005. Person identification from biological motion: effects of structural and kinematic cues. *Perception & Psychophysics*, 67, 4, 667–675. ISSN: 1532-5962. DOI: <https://doi.org/10.3758/BF03193523>.
- [50] Rick van der Zwan, C MacHatch, Desiree Kozlowski, NF Troje, O Blanke, and Anna Brooks. 2009. Gender bending: auditory cues affect visual judgements of gender in biological motion displays. *Experimental Brain Research*, 198, 2, 373–382. DOI: <https://doi.org/10.1007/s00221-009-1800-y>.
- [51] Changsheng Wan, Li Wang, and Vir V. Phoha. 2018. A survey on gait recognition. *ACM Comput. Surv.*, 51, 5. ISSN: 0360-0300. DOI: <https://doi.org/10.1145/3230633>.
- [52] Hao Wang, Zhengquan Xu, Shan Jia, Ying Xia, and Xu Zhang. 2021. Why current differential privacy schemes are inapplicable for correlated data publishing? *World Wide Web*, 24, 1, 1–23. DOI: <https://doi.org/10.1007/s11280-020-00825-8>.
- [53] Galit Yovel and Alice J O’Toole. 2016. Recognizing people in motion. *Trends in cognitive sciences*, 20, 5, 383–395. ISSN: 1364-6613. DOI: <https://doi.org/10.1016/j.tics.2016.02.005>.
- [54] Yuhan Zhou, Robbin Romijnders, Clint Hansen, Jos van Campen, Walter Maetzler, Tibor Hortobágyi, and Claudine JC Lamoth. 2020. The detection of age groups by dynamic gait outcomes using machine learning approaches. *Scientific reports*, 10, 1, 1–12. ISSN: 2045-2322. DOI: <https://doi.org/10.1038/s41598-020-61423-2>.