

Article

# Learning Gait Representations with Noisy Multi-Task Learning

Adrian Cosma  and Emilian Radoi \* Faculty of Automatic Control and Computer Science, University Politehnica of Bucharest,  
006042 Bucharest, Romania

\* Correspondence: emilian.radoi@upb.ro

**Abstract:** Gait analysis is proven to be a reliable way to perform person identification without relying on subject cooperation. Walking is a biometric that does not significantly change in short periods of time and can be regarded as unique to each person. So far, the study of gait analysis focused mostly on identification and demographics estimation, without considering many of the pedestrian attributes that appearance-based methods rely on. In this work, alongside gait-based person identification, we explore pedestrian attribute identification solely from movement patterns. We propose DenseGait, the largest dataset for pretraining gait analysis systems containing 217 K anonymized tracklets, annotated automatically with 42 appearance attributes. DenseGait is constructed by automatically processing video streams and offers the full array of gait covariates present in the real world. We make the dataset available to the research community. Additionally, we propose GaitFormer, a transformer-based model that after pretraining in a multi-task fashion on DenseGait, achieves 92.5% accuracy on CASIA-B and 85.33% on FVG, without utilizing any manually annotated data. This corresponds to a +14.2% and +9.67% accuracy increase compared to similar methods. Moreover, GaitFormer is able to accurately identify gender information and a multitude of appearance attributes utilizing only movement patterns. The code to reproduce the experiments is made publicly.

**Keywords:** gait recognition; self-supervised learning; pose estimation; multi-task learning; weakly-supervised learning



**Citation:** Cosma, A.; Radoi, E. Learning Gait Representations with Noisy Multi-Task Learning. *Sensors* **2022**, *22*, 6803. <https://doi.org/10.3390/s22186803>

Academic Editors: Euntai Kim, Sangyoun Lee and Kang Ryoung Park

Received: 12 August 2022  
Accepted: 6 September 2022  
Published: 8 September 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Technologies relying on facial and pedestrian analysis play a crucial role in intelligent video surveillance and security systems. Facial and pedestrian analysis systems have become the norm in video intelligence, such systems being deployed ubiquitously. However, appearance-based pedestrian re-identification [1] and facial recognition models [2] invariably suffer from extrinsic factors related to camera viewpoint and resolution, and to the change in a person's appearance such as different clothing, hairstyles and accessories. Moreover, due to the proliferation of privacy laws such as GDPR, it is increasingly difficult to deploy appearance-based solutions for video-intelligence. Human movement is highly correlated with many internal and external aspects of a particular individual including age, gender, body mass index, clothing, carrying conditions, emotions and personality [3]. The manner of walking is unique to each person, it does not significantly change in short periods of time [4] and cannot be easily faked to impersonate another person [5]. Gait analysis has gained significant attention in recent years [6,7], due to solving many of the problems of appearance-based technologies without relying on the direct cooperation of subjects. However, compared to appearance-based methods, gait analysis is intrinsically harder to perform with reliable accuracy, due to the influence of many confounding factors that affect the manner of walking. This problem is tackled in literature in two major ways, either by building specialized neural architectures that are invariant to walking variations [8–10], or by creating large-scale and diverse datasets for training [11–15].

One of the first attempts of building a large-scale gait recognition dataset is OU-ISIR [14], which is comprised of 10,307 identities that walk in a straight line for a short duration

of time. Such a dataset is severely limited by its lack of walking variability, having only viewpoint change as a confounding factor. Building sufficiently large datasets that account for all the walking variations imply an immense annotation effort. For example, the GREW benchmark [12] for gait-based identification, reportedly took 3 months of continuous manual annotation by 20 workers. In contrast, automatic, weakly annotated datasets are much easier to gather by leveraging existing state-of-the-art models—UWG [11], a comparatively large dataset of individual walking tracklets proved to be a promising new direction in the field. Increasing the dataset size is indeed correlated with performance on downstream gait recognition benchmarks [11], even though no manual annotations are provided. One limitation of these datasets is that they are annotated with attributes per individual only sparsely, and not addressing the problem of pedestrian attribute identification (PAI), currently performed only through appearance-based methods [16–18]. Walking pedestrians are often annotated only with their gender, age, and camera viewpoint [8,12,14,15]. Even though gait-based demographic identification is a viable method for pedestrian analysis [19], it is also severely limited by the lack of data. Also, many attributes from PAI networks such as gender, age and body type have a definite impact on walking patterns [20–22], and we posit that they can be identified with a reasonable degree of accuracy using only movement patterns and not utilizing appearance information.

We propose DenseGait, the largest gait dataset for pretraining to date, containing 217 K anonymized tracklets in the form of skeleton sequences, automatically gathered by processing real-world surveillance streams through state-of-the-art models for pose estimation and pose tracking. An ensemble of PAI networks was used to densely annotate each skeleton sequence with 42 appearance attributes such as their gender, age group, body fat, camera viewpoint, clothing information and apparent action. The purpose of DenseGait is to be used for pretraining networks for gait recognition and attribute identification, it is not suitable for evaluation since it is annotated automatically and does not contain manual, ground-truth labels. DenseGait contains walking individuals in real scenarios, it is markerless, non-treadmill, and avoids unnatural and constrictive laboratory conditions, which have been shown to affect gait [23]. It practically contains the full array of factors that are present in real world gait patterns.

The dataset is fully anonymized, and any information pertaining to individual identities is removed, such as the time, location and source of the video stream, and the appearance and height information of the person. DenseGait is a gait analysis dataset primarily intended for pretraining neural models—using it to explicitly identify the individuals within it is highly unfeasible, requiring extensive external information about the individuals, such as personal identifying information (i.e., their name or ID) and a baseline gait pattern. According to GDPR (<https://eur-lex.europa.eu/eli/reg/2016/679/oj>, accessed on 1 July 2022) legislation, data used for research purposes can be used if anonymized. Moreover, anonymized data does not conform to the rigors of personal data and can be processed without explicit consent. Nevertheless, *any attempt to use of DenseGait to explicitly identify individuals present in it is highly discouraged.*

We chose to utilize only skeleton sequences for gait analysis, as current appearance-based methods that rely on silhouettes are not privacy preserving, potentially allowing for identification based only on the person's appearance, rather than their movement [24]. Skeleton sequences encode only the movement of the person, abstracting away any visual queues regarding identity and attributes. Moreover, skeleton-based solutions have the potential to generalize across tasks such as action recognition, allowing for a flexible and extensible computation.

DenseGait, compared to other similar datasets [11], contains  $10\times$  more sequences and is automatically annotated with 42 appearance attributes through a pretrained PAI ensemble (Table 1). In total, 60 h of video streams were processed, having a cumulative walking duration of pedestrians of 410 h. We release the dataset under open credentialized access, for research purposes only, under CC-BY-NC-ND-4.0 (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>, accessed on 1 July 2022) License.

**Table 1.** List of attributes extracted by each network in the PAI ensemble. Each network is trained on a different dataset, with a separate set of attributes. After coalescing similar attributes and eliminating appearance-only attributes, we obtain 42 appearance attributes.

| PA100k  | PETA   | RAP  |
|---|--|--|
| Female, AgeOver60, Age18–60, Front, Side, Back, Hat, Glasses, HandBag, ShoulderBag, Backpack, HoldObjectsInFront, ShortSleeve, LongSleeve, UpperStride, UpperLogo, UpperPlaid, UpperSplice, LowerStripe, LowerPattern, LongCoat, Trousers, Shorts, Skirt & Dress, Boots | Age16–30, Age31–45, Age46–60, AgeAbove61, Backpack, CarryingOther, Casual lower, Casual upper, Formal lower, Formal upper, Hat, Jacket, Jeans, Leather-Shoes, Logo, LongHair, Male, Messenger Bag, Muffler, No accessory, No carrying, Plaid, PlasticBags, Sandals, Shoes, Shorts, Short Sleeve, Skirt, Sneaker, Stripes, Sunglasses, Trousers, TShirt, UpperOther, V-Neck | Female, AgeLess16, Age17–30, Age31–45, BodyFat, BodyNormal, BodyThin, Customer, Clerk, Bald-Head, LongHair, BlackHair, Hat, Glasses, Muffler, Shirt, Sweater, Vest, TShirt, Cotton, Jacket, Suit-Up, Tight, ShortSleeve, LongTrousers, Skirt, ShortSkirt, Dress, Jeans, TightTrousers, LeatherShoes, SportShoes, Boots, ClothShoes, CasualShoes, Backpack, SSBag, Hand-Bag, Box, PlasticBags, PaperBag, HandTrunk, OtherAttachment, Calling, Talking, Gathering, Holding, Pusing, Pulling, CarryingbyArm, CarryingbyHand |

We also propose GaitFormer, a multi-task transformer-based architecture [25] that is pretrained on DenseGait in a self-supervised manner, being able to perform exceptionally well in zero-shot gait recognition scenarios on benchmark datasets, achieving 92.5% identification accuracy from direct transfer on the popular CASIA-B dataset, without using any manually annotated data. Moreover, it obtains good results on demographic and pedestrian attribute identification from walking patterns, with no manual annotations. GaitFormer represents the first use of a plain transformer encoder architecture in gait skeleton sequence processing, without relying on hand-crafted architectural modifications as in the case of graph neural networks [26,27].

This paper makes the following contributions:

1. We release DenseGait, the largest dataset of skeleton walking sequences, densely annotated with appearance information, for use in pretraining neural architectures that can be further fine-tuned on specific gait analysis tasks. The dataset can be found at <https://bit.ly/3SLO8RW>, under open credentialized access, for research purposes only.
2. We propose GaitFormer, a multi-task transformer that is pretrained on the DenseGait dataset and achieves exceptional results in zero-shot gait recognition scenarios on benchmark datasets, achieving 92.52% accuracy on CASIA-B and 85.33% on FVG, without training on any manually annotated data (+14.2% and +9.67% increase compared to similar methods [11]). The code is made publicly available at: <https://github.com/cosmaadrian/gaitformer>
3. We explore the performance of GaitFormer on other gait analysis tasks, such as gait-based gender estimation and attribute identification.

## 2. Related Work

### 2.1. Gait Analysis

Video gait analysis encompasses research efforts dedicated to automatically estimate and predict various aspects of a walking person. Research has been mostly dedicated into gait-based person recognition, with many benchmark datasets [8,11,12,14,28–31] available for training and testing models. Moreover, there have been improvements in areas such as estimating demographics information [19,32], emotion detection [33] and ethnicity estimation [34] from only movement patterns. Sepas-Moghaddam and Etemad [35] proposed a taxonomy to organize the existing works in the field of gait recognition. In this work, we focus mainly on body representation, as we made a deliberate choice of providing DenseGait

with only movement information for anonymization. Broadly, works in gait analysis can be divided into two major approaches in terms of body representation: silhouette-based and skeleton-based.

### 2.1.1. Silhouette-Based Solutions

Silhouette-based approaches make use of silhouettes of walking individuals estimated either through background subtraction methods or through instance segmentation and tracking. Silhouettes are used in various forms, either in a condensed representation [36–38], or as a sequences, as it is the norm in more modern methods [8,39–41]. Most notably, GaitSet [39] processes the silhouettes as a set, as opposed to preserving the temporal information present in a sequence. As such, the authors can include silhouettes from multiple videos of the same walking subjects, achieving good invariance to walking variations. GaitPart [40] processes the temporal variation of each individual body part separately in a Micro-motion Capture Module (MCM), taking inspiration from model-based approaches. Each body part exhibits different visual queues and temporal variation and the authors propose to combine the each feature part to construct the final gait representation. Recently, Lin et al. [41], advance the construction of neural architectures for processing silhouette sequences by proposing a Global-Local Feature Extractor (GLFE), which obtains good results on benchmark datasets. Zhang et al. [8] propose GaitNet, a model which directly makes use of the appearance of the individual and is able to output invariant feature representations for gait recognition. Moreover, they also propose FVG, a dataset with 226 individuals, only from the front-view angle, one of the more challenging angles in gait analysis, due to the lack of perceived variation in limb movements.

### 2.1.2. Skeleton-Based Solutions

Skeleton-based approaches, on the other hand, avoid making use of appearance information in the form of silhouettes, and instead focus on the moving anatomical skeleton of the person, effectively processing only movement patterns. Approaches typically imply processing walking sequences with a pose estimation [42] model, and processing the resulting skeletons with a neural network, either by adapting conventional CNN modules [43], or with an LSTM [44,45]. More modern approaches make use of graph neural networks to model the relationships between human joints [46,47]. Liao et al. [44] make use of a combined CNN and LSTM architecture to model 2D skeleton sequences. A later improvement makes use of 3D skeletons [45] to further improve results. Li et al. [46] propose a graph-based convolutional architecture to process skeleton sequences, and a Joints Relationship Pyramid Mapping to map spatio-temporal gait features into a discriminative feature space. Li and Zhao [47] propose CycleGait, a graph-based approach that incorporates multiple walking paces in the augmentation procedure and obtains robust results in gait recognition on CASIA-B. In contrast to these approaches, we opted to take a data-driven approach, instead of an algorithmic approach, and use a standard transformer architecture and pretrain it on a large amount of weakly-labelled data. Recently, Cosma and Radoi [11] proposed an approach called WildGait to skeleton-based gait recognition, in which they automatically mine surveillance streams and pretrain a ST-GCN [26] model in a self-supervised manner. Through fine-tuning, good results are obtained in recognition on CASIA-B and FVG. Similarly to WildGait, we also process publicly available surveillance streams, but increase the DenseGait dataset size by an order of magnitude. Moreover, we densely annotate each skeleton sequence with 42 appearance attributes for use in zero-shot attribute identification scenarios.

However, model-based approaches still lag behind methods utilizing appearance (i.e., silhouettes). This is most likely due to the imperfect extraction of skeletons by modern pose estimators, which struggle to accurately detect fine-grained movements at a distance. Moreover, using appearance-based methods is fundamentally easier, since a single silhouette can contain identifying information about a subject. For instance Xu et al. [48] obtained reasonable results for gait recognition using a single silhouette, which cannot be considered

gait, as no temporal movement is being processed at all. This implies that recognition is performed through “shortcuts” in the form of appearance features (i.e., body composition, height, haircut, side-profile etc). For this reason, a more privacy-aware approach is to process only movement patterns, which constitutes the motivation for releasing DenseGait with only anonymized skeleton sequences, and disregarding silhouettes.

## 2.2. Transformers and Self-Supervised Learning

In recent years, there has been an insurgence of research in the area of self-supervised learning, mostly due to the extremely high performance obtained in natural language processing with models such as BERT [49] and GPT [50]. Self-supervised learning presumes training models using aspects of the data itself as a supervisory signal. While initial efforts in computer vision relied on creating artificial pretext tasks [51–53], the field is moving towards contrastive-based approaches [29,54,55]. Methods such as SimCLR [29], Barlow Twins [54] and Dino [55] obtaining almost similar performance to direct supervision. Moreover, the transformer has proven to be a flexible architecture, capable of handling a multitude of modalities such as text [49], images [56], video [57], speech [58], and highly benefit from large-scale pretraining [59]. Taking inspiration from related efforts to process non-textual data with transformers [56], we construct GaitFormer by processing flattened skeletons as input “tokens”. In this manner, any human bias related to hand-crafted graph relationships between the body joints is eliminated. Moreover, as opposed to graph networks such as ST-GCN [27], training a similarly large transformer encoder make more efficient use of computational resources, significantly reducing training time.

## 3. Method

### 3.1. Dataset Construction

For building the DenseGait dataset, we made use of public video streams (e.g., street cams), and processed them with AlphaPose [42], a modern, state-of-the-art multi-person pose estimation model. AlphaPose’s raw output is comprised of skeletons with  $(x, y, c)$  coordinates for each of the 18 joints of the COCO skeleton format Lin et al. [60], corresponding to 2D coordinates in the image plane and a prediction confidence score. We performed intra-camera tracking for each skeleton with on SortOH [61]. SortOH is based on the SORT [62] algorithm, which relies only on coordinate information and not on appearance information. As opposed to DeepSORT [63] which makes use of person re-identification models, SortOH is only using coordinates and bounding box size for faster computation time while having comparably similar performance. SortOH ensures that tracking is not significantly affected by occlusions.

To ensure that the skeleton sequences can be properly processed by a deep learning model, we performed extensive data cleaning. We have filtered low confidence skeletons by computing the average confidence of each of the 18 joints, and in each sequence, skeletons with an average confidence of less than 0.5 were removed. Furthermore, skeletons with feet confidence less than 0.4 were removed. This step guarantees that the feet are visible and confidently detected—leg movement is one of the most important signals for gait analysis. In our processing, we chose a period length  $T$  of 48 frames, which corresponds to approximately 2 full gait cycles on average [64]. Surveillance streams do not have the same frame rate between them, which makes the sequences have different paces and durations. As such, we filtered short tracklets which have a duration of less than  $\frac{T * fps}{24}$ . We consider 24 FPS to be real-time video speed, and each video was processed according to its own frame rate. Moreover, skeletons are linearly interpolated such that the pace and duration is unified across video streams.

Similar to [11], we further normalized each skeleton by centering at the pelvis coordinates  $(x_{pelvis}, y_{pelvis})$  and scaling vertically by the distance between the head and the hips  $(y_{neck} - y_{pelvis})$  and horizontally by the distance between the shoulders  $((x_{R.shoulder} - x_{L.shoulder}))$ . This procedure is detailed in Equations (1) and (2). The normalization procedure aligns the skeleton sequences in a similar manner to the alignment step in

face recognition pipelines [65]. This step eliminated the height and body type information about the subject, ensuring that the person cannot be directly identified.

$$x_{joint} = \frac{x_{joint} - x_{pelvis}}{|x_{R.shoulder} - x_{L.shoulder}|} \quad (1)$$

$$y_{joint} = \frac{y_{joint} - y_{pelvis}}{|y_{neck} - y_{pelvis}|} \quad (2)$$

However, body type information should be preserved through the analysis of the walking patterns. Moreover, normalization obscures the human position in the frame, to prevent identification of the source video stream.

Finally, we filtered standing/non-walking skeletons in each sequence by computing the average movement speed of the legs, which is indicative of the action the person is performing. As such, if the average leg speed is less than 0.0015 and higher than 0.09, the sequence was removed. The thresholds were determined through manual inspection of the sequences. This eliminated both standing skeleton sequences as well as sequences with erratic leg movement, which is most probably due to poor pose estimation output in that case.

DenseGait is fully anonymized. Any information regarding the identity of particular individuals in the dataset is eliminated, including appearance information (by keeping only movement information in the form of skeleton sequences), height and body proportions (through normalization), and the time, location, and source of the video stream. Identifying individuals in DenseGait is highly unfeasible, as it requires external information (i.e., name, email, ID, etc.) and specific collection of gait patterns.

The final dataset contains 217 K anonymized tracklets, with a combined length of 410 h. DenseGait is currently the largest dataset of skeleton sequences for use in pretraining gait analysis models. Table 2 showcases a comparison between DenseGait and other popular gait recognition datasets. Since the skeleton sequences are collected automatically through pose tracking, it is impossible to quantify exactly the number of different identities in the dataset, as, in some cases, tracking might be lost due to occlusions. However, DenseGait contains a significantly larger number of tracklets compared to other available datasets while also being automatically densely annotated with 42 appearance attributes. In the case of UWG [11] and DenseGait, the datasets do not contain explicit covariates for each identity, but rather covariates in terms of viewing angle, carrying conditions, clothing change, and apparent action are present across the tracklet duration, similar to GREW [12].

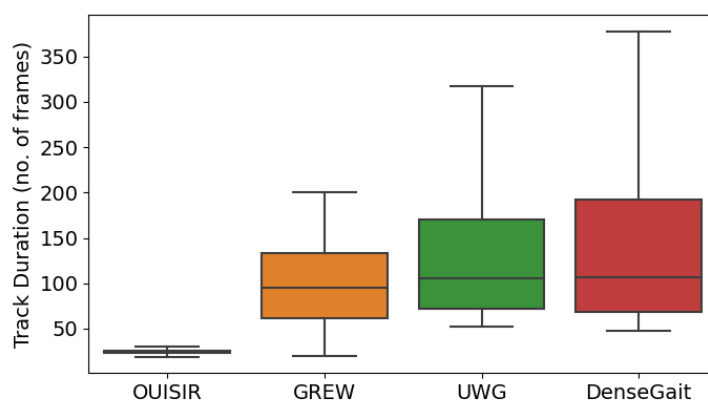
Similarly to UWG [11], DenseGait does not contain multiple walks per person, rather each tracklet is considered a unique identity. Compared to other large-scale datasets, DenseGait tracks individuals for a longer duration, which makes it suitable for use in self-supervised pretraining, as longer tracked walking usually contains more variability for a single person. Figure 1 shows boxplots with a five-number summary descriptive statistics for the distribution of track durations in each dataset. DenseGait has a mean tracklet duration of 162 frames, which is significantly larger (z-test  $p < 0.0001$ ) compared to other datasets: CASIA-B [15]—83 frames, FVG [8]—97 frames, GREW [12]—98 frames, UWG [11]—136 frames). Due to potential loss of tracking information, the dataset is noisy, and can be used only for self-supervised pretraining.

### 3.2. Annotations with Appearance Attributes

Appearance attributes are essential for pretraining for tasks such as gender estimation [19], age estimation [66] and pedestrian attribute identification [16–18]. To ensure that the dataset is densely annotated with appearance attributes, we made use of an ensemble of pretrained PAI networks, each trained on different popular PAI datasets. Specifically, we employed three InceptionV3 [67] networks trained on RAP [68], PETA [69] and PA100k [16], respectively. Figure 2 showcases the annotation procedure.

**Table 2.** Comparison of popular datasets for gait recognition. DenseGait is an order of magnitude larger, has more identities in terms of skeleton sequences (highlighted in **bold**), and each sequence is annotated with 42 appearance attributes. \* Approximate number given by pose tracker. † Implicit covariates across tracking duration.

| Dataset                 | # IDs            | Sequences      | Covariates | Views | Env.    |
|-------------------------|------------------|----------------|------------|-------|---------|
| USF HumanID [31]        | 122              | 1870           | Y          | 2     | Outdoor |
| TUM-GAID [28]           | 305              | 3370           | Y          | 1     | Outdoor |
| FVG [8]                 | 226              | 2857           | Y          | 1     | Outdoor |
| CASIA-B [15]            | 124              | 13,640         | Y          | 11    | Indoor  |
| OU-ISIR [14]            | 10,307           | 144,298        | N          | 14    | Indoor  |
| GREW [12]               | 26,000           | 128,000        | Y          | -     | Outdoor |
| UWG [11]                | 38,502 *         | 38,502         | Y †        | -     | Outdoor |
| <b>DenseGait (ours)</b> | <b>217,954 *</b> | <b>217,954</b> | Y †        | -     | Outdoor |



**Figure 1.** Comparison between existing large-scale skeleton gait databases and DenseGait in terms of distributions of tracklet duration. DenseGait is an order of magnitude larger than the next largest skeleton database, while having a longer average duration (136 frames UWG vs 162 frames DenseGait).

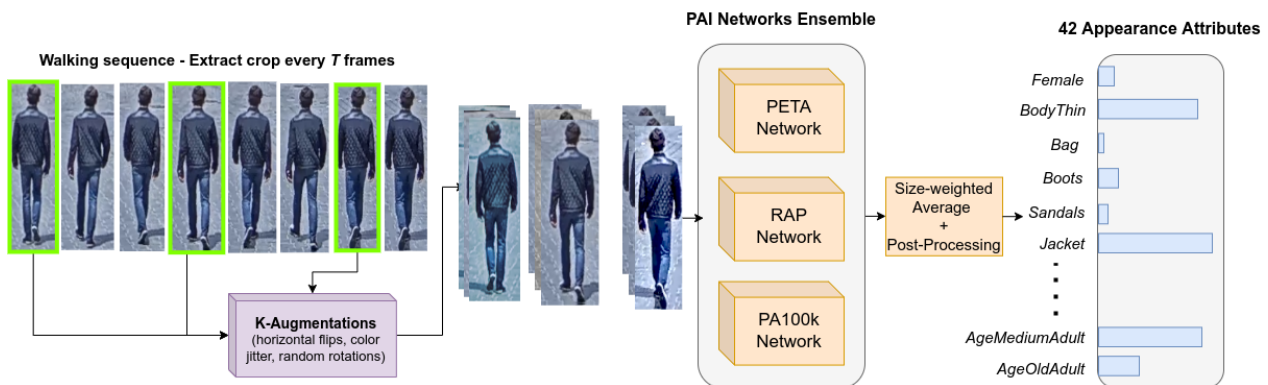
Since each dataset has a different set of pedestrian attributes, we averaged similar classes (e.g., *AgeLess16* and *AgeLess18* into *AgeChild*), coalesced similar classes (e.g., *Formal* and *Suit-Up* into *FormalWear*) and removed attributes that cannot evidently be estimated from movement patterns (e.g., *BaldHead*, *Hat*, *V-Neck*, *Glasses*, *Plaid* etc.).

For a particular sequence, we take the cropped image of the pedestrian at every  $T$  frames (where  $T$  is the period length), and randomly augment it  $k = 4$  times (e.g., random horizontal flips, color jitter and small random rotation). For each crop, each augmented version is then processed by a PAI network and the results are averaged such that the output is robust to noise [70]. Finally, to have a unified prediction for the walking sequence, results are averaged according to the size of the bounding box relative to the image, similar to Catruna et al. [19]. Predictions on larger crops have a higher weight, with the assumption that the pedestrian appearance is more clearly distinguishable when closer to the camera.

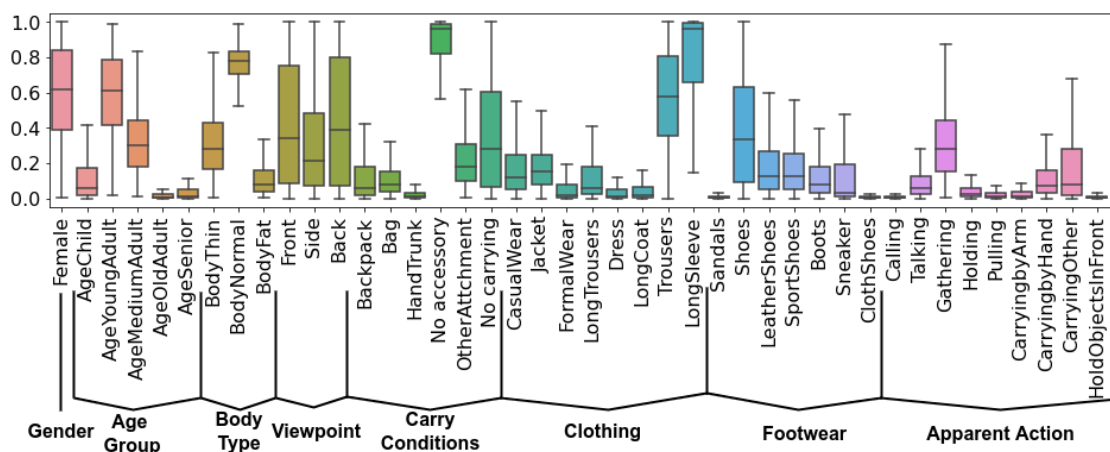
Figure 3 showcases the final list of attributes, and their distribution across the dataset. We have a total of 42 attributes, split into 8 groups: *Gender*, *Age Group*, *Body Type*, *Viewpoint*, *Carry Conditions*, *Clothing*, *Footwear* and *Apparent Action*. For the final annotations, we chose to keep the soft-labels and not round them, as utilizing soft-labels for model training was shown to be a more robust approach when dealing with noisy data [70].

Figure 4 showcases selected examples of attribute predictions from the PAI ensemble. Since surveillance cameras usually have low resolution and the subject might be far away from the camera, some pedestrian crops are blurry and might affect prediction by the PAI ensemble. For gender, age group, body composition and viewpoint, the models are confidently identifying these attributes. However, for specific pieces of clothing (i.e., footwear: *Sandals/LeatherShoes*), predictions are not always reliable, due to the low

resolution of some of the crops, but the errors are negligible when taking into account the scale of the dataset.



**Figure 2.** Overview of the automatic annotation procedure for the 42 appearance attributes. To robustly annotate attributes, an ensemble of pretrained networks is used in conjunction with multiple augmentations of the same crop. Predictions across the sequence are averaged according to their bounding-box area.



**Figure 3.** Distribution of the 42 appearance attributes in DenseGait. The dataset is annotated in a fine-grained manner with attributes ranging from internal aspects of the person (Gender, Age Group, Body Type) to appearance only labels (Clothing, Footwear).



**Figure 4.** Qualitative examples for selected attributes from the PAI ensemble. The networks correctly identify gender, age group and viewpoint. However, in some cases, clothing and, more specifically, footwear are more difficult to estimate in low resolution scenarios.

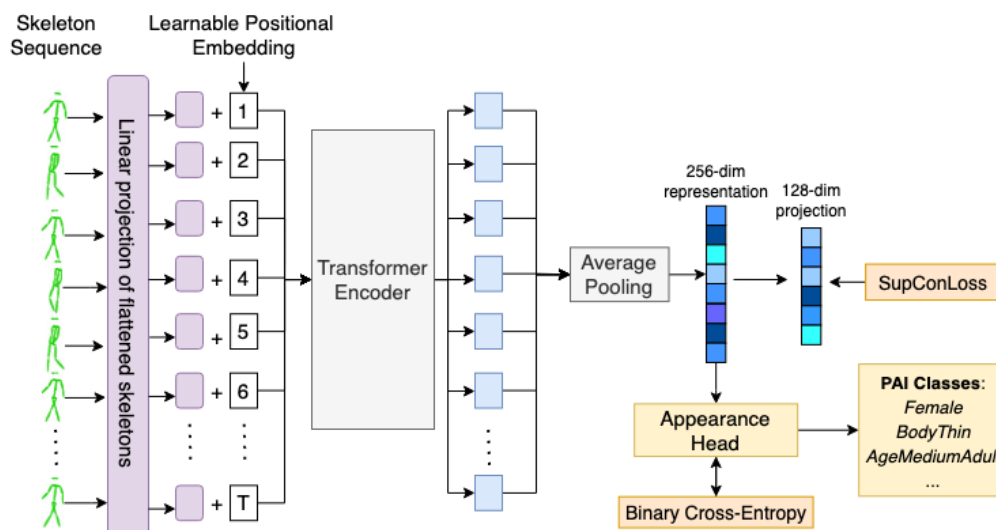


### 3.3. Description of Model Architecture

For pretraining on the DenseGait dataset for the tasks of gait-based recognition and attribute identification, we chose to adapt the popular transformer encoder architecture [25] to handle skeleton sequences. Initially, transformers were immensely successful in handling sequential data in the form of text, effectively replacing LSTM [71] networks, the de facto approach for these problems. However, lately, transformers have been used in a variety of problems, being able to handle images [56], video [57] and multi-modal data [72]. Moreover, transformer architectures in particular highly benefit from large-scale, self-supervised pretraining [49,50,55], allowing models to be effectively fine-tuned on more specific datasets with small amounts of annotated data.

To handle skeleton sequences, we abstain from making any hand-crafted architectural modifications, as in the case of Plizzari et al. [27], which uses a hybrid approach by combining graph computation on the skeleton and using multi-head attention on the extracted features. Instead, we take inspiration from ViT [56], which processes images as a sequence of flattened patches that are fed into a standard transformer encoder network. Figure 5 showcases the training procedure for GaitFormer in the multi-task training regime. Each skeleton is flattened into a 54 dimensional vector and is linearly projected with a standard learnable feed-forward layer into a 256 dimensional space. Each skeleton projection is then fed into a transformer encoder network. We opted for learnable positional embedding that is added to each projection instead of concatenated, to avoid increasing the dimensionality. After the transformer encoder, representations for each skeleton are averaged, and a final linear feed-forward layer of 256 elements is used as the final embedding. Further, as described in SimCLR [29], we used an additional 128-dimension linear layer for training with a supervised contrastive objective [73]. Additionally, a linear layer is used as appearance head to estimate the pedestrian attributes that is trained using a standard binary-crossentropy loss.

We used three different model sizes for the transformer encoder in our experiments, with 4 encoder layers (*SM*), 8 encoder layers (*MD*) and 12 encoder layers (*XL*). In all types of architectures, 8 attention heads were used, and the internal feed-forward dimensionality was 256 [25].



**Figure 5.** Overview of GaitFormer (Multi-Task) training procedure. Flattened skeletons are linearly projected using a standard feed-forward layer and fed into a transformer encoder. The vectorized representations are average pooled and the resulting 256-dimensional vector is used for estimating the identity and to estimate the 42 appearance attributes through the “Appearance Head”. The contrastive objective (SupConLoss) is applied to a lower 128-dimensional linear projection, similar to the approach in SimCLR [29].

### 3.4. Training Details

For training on DenseGait, we chose to use contrastive learning [29] as a supervisory signal. By design, contrastive methods work by attracting representations belonging to the same class, while simultaneously repelling samples from different classes. This paradigm is identical to the objective for recognition problems, which constitutes one of the main tasks in gait analysis. Specifically, we used SupConLoss [73], with a temperature of  $\tau = 0.001$ , alongside a two-view sampler for each skeleton in the batch. SupConLoss assumes a multi-viewed batch, with multiple augmentations for the same sample. Each view of a skeleton sequence is randomly augmented by the standard suite of augmentations for this data modality: random sequence crops of fixed length of  $T = 48$ , random flips with 50% probability, random paces [53], and random gaussian noise added to joints coordinates. Let  $i \in I \equiv \{1 \dots 2N\}$  be the index of an arbitrary augmented sample. SupConLoss is defined as:

$$\mathcal{L}^{sup} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \in A(i)} \exp(z_i \cdot z_a / \tau)} \quad (3)$$

In Equation (3),  $z_i = Enc(\tilde{x}_i)$  denotes the embedding of a skeleton sequence  $x_i$ , “ $\cdot$ ” denotes the dot product operation and  $A(i) \equiv I \setminus \{i\}$ . Moreover,  $P(i) \equiv \{p \in A(i) : \tilde{y}_p = \tilde{y}_i\}$  is the set of indices of all positives in the multi-viewed batch distinct from  $i$ . In our case, the positive pairs are constructed by two different augmentations of the same skeleton sequence. The variability of the two augmentations is higher if the skeleton is tracked for a longer duration of time, as the walking individual might change direction.

As suggested in Chen et al. [29], the supervisory signal given by SupConLoss is applied to a lower dimensional embedding (128 dimensions) to avoid the curse of dimensionality.

For predicting appearance attributes, which is a multi-label problem, we used a standard binary-crossentropy loss between each appearance label ( $p_i$ ) and its corresponding prediction ( $y_i$ ) (Equation (4)). As previously mentioned, we keep the soft labels as a supervisory signal, to prevent the network from overfitting and be more robust to noisy or incorrect labels [74]. Moreover, since learning appearance labels can be regarded as a knowledge distillation problem between the PAI ensemble and the transformer network, soft labels help improve the distillation process [75].

$$\mathcal{L}_{appearance} = -(y_i \log(p_i) + (1 - y_i) \log(1 - p_i)) \quad (4)$$

In multi-task (MT) training scenarios, we used a combination of the two losses, with a weight penalty of  $\lambda = 0.5$  on the appearance loss  $\mathcal{L}_{appearance}$ . We chose  $\lambda = 0.5$  empirically, such that the two losses have similar magnitudes. The final loss function is defined as:

$$\mathcal{L}_{final} = \mathcal{L}_{SupCon} + \lambda \mathcal{L}_{appearance} \quad (5)$$

In plain contrastive training scenarios, we employ only the SupConLoss, without predicting attributes (i.e.,  $\mathcal{L}_{final} = \mathcal{L}_{SupCon}$ ).

The motivation for pre-training the network in a multi-task setting is that the network not only learns to cluster walking sequences by their identity, but also to take appearance attributes into account. For instance, predicting the gender and age, even if they are not completely reliable, could prove useful for gait recognition, as demographics can be considered soft-biometrics, allowing the network to automatically filter identities by these attributes. On the other hand, in contrastive-only scenario, the network is under a classical self-supervised regime.

We used a batch size of 1024 across our experiments, with a cyclical learning rate [76] ranging from 0.0001 and 0.001 across 20 epochs. We trained all models for 400 epochs.

## 4. Experiments and Results

This section explores the performance of GaitFormer on gait-based recognition, gender identification and pedestrian attribute identification. We are primarily interested in evalu-

ating the model in scenarios with low amounts of annotated data and we opted to use the two popular benchmark datasets originally constructed for gait recognition: CASIA-B [15] and FVG [8]. For gender estimation, we manually annotated the gender information for each identity in the two datasets and constructed CASIA-gender and FVG-gender. We briefly describe each dataset below.

We chose CASIA-B to compare with other skeleton-based gait recognition models, since it is one of the most popular gait recognition datasets in literature. It contains 124 subjects walking indoors in a straight line, captured with 11 synchronized cameras with three walking variations—normal walking (NM), clothing change (CL) and carry conditions (BG). According to Yu et al. [15], the first 62 subjects are used for training and the rest for evaluation. CASIA-gender consists of manually annotated the subjects in CASIA-B with gender information, having a split of 92 males and 32 females. We maintain the training and validation splits from the recognition task, using the first 62 subjects for training (44 males and 18 females) and the rest for validation (48 males and 14 females). We use FVG to evaluate the robustness of GaitFormer, as it contains different covariates than CASIA-B such as varying degrees of walking speed, the passage of time and cluttered background. Moreover, FVG only contains walks from the front-view angle, which is more difficult for gait processing due to lower perceived limb variation. According to Zhang et al. [8], from the 226 identities present in FVG, the first 136 are used for training and the rest for testing. Similarly, FVG-gender contains manual annotations with gender information, obtaining 149 males and 77 females. We maintain the training and validation splits from the recognition task, utilizing the first 136 individuals for training (83 males and 53 females) and the rest for validation (66 males and 24 females).

#### 4.1. Recognition

We initially trained GaitFormer under two regimes: (i) contrastive only and (ii) multi-task (MT), which implies training with SupConLoss [73] on the tracklet ID while simultaneously estimating the appearance attributes (Figure 5). We experiment with three models sizes: SM—4 encoder layers (2.24M parameters), MD—8 encoder layers (4.35M parameters) and XL—12 encoder layers (6.46M parameters).

We pretrain GaitFormer on the DenseGait dataset under the mentioned conditions and directly evaluate recognition performance in terms of accuracy on CASIA-B and FVG, without fine-tuning. In all experiments we perform a deterministic crop in the middle of the skeleton sequences of  $T = 48$  frames, and use no test-time augmentations. For each cropped skeleton sequence, features are extracted using the 256-dimensional representation and are normalized with the  $l_2$  norm. In Table 3 we present results on the walking variations for each model size and training regime. For CASIA-B, we show mean accuracy where the gallery set contains all viewpoints except the probe angle, in the three evaluation scenarios: normal walking (NM), change in clothing (CL) and carry bag (CB). For FVG, we show accuracy results based on the evaluation protocols mentioned by Zhang et al. [8], corresponding to different walking scenarios (walk speed (WS), change in clothing (CL), carrying bag (CB), cluttered background (CBG) and ALL). Results show that unsupervised pretraining on DenseGait is a viable way to perform gait recognition, achieving an accuracy of 92.52% on CASIA-B and 85.33% on FVG, without any manually annotated data available. Notably, multi-task learning on appearance attributes provides a consistent positive gap in the downstream performance.

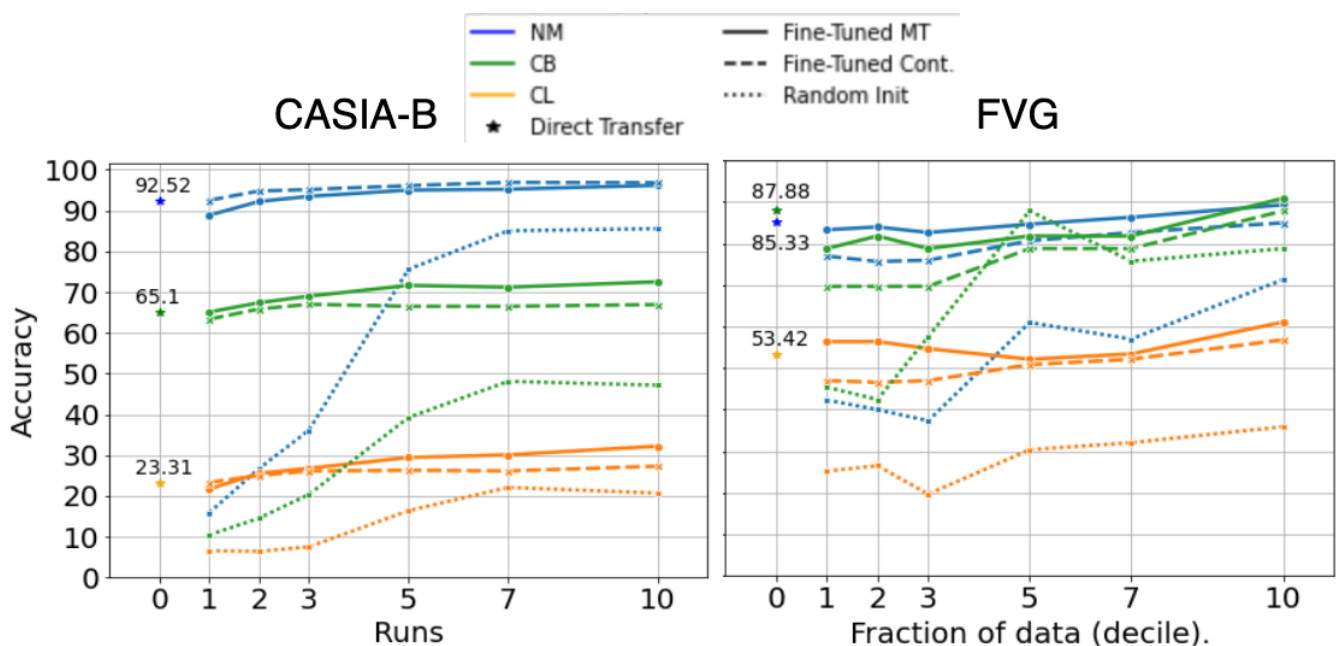
Model size in terms of number of layers does not seem to considerably affect performance on benchmark datasets. GaitFormerMD (8 layers) fairs consistently better than GaitFormerXL (12 layers), while being similarly close to GaitFormerSM (4 layers).

Figure 6 compares GaitFormerMD pretrained on DenseGait in the two training regimes (contrastive only—Cont. and Multi-Task—MT) and GaitFormerMD randomly initialized. The networks were fine-tuned on progressively larger samples of the corresponding datasets: for CASIA-B, we sampled multiple runs for the same identity (from 1 to 10 runs per ID), and for FVG, we randomly sampled a percentage of runs per each identity. Mod-

els were fine-tuned using Layer-wise Learning Rate Decay (LLRD) [77], which implies a higher learning rate for top layers and a progressively lower learning rate for bottom layers. The learning rate was decreased linearly from 0.0001 to 0, across 200 epochs. The results show that unsupervised pretraining has a substantial effect on downstream performance especially in low data scenarios (direct transfer and 10% of available data). Moreover, pretraining the model in Multi-Task learning regime, in which the network was tasked to estimate appearance attributes from movement alongside with the identity, provides a consistent increase in performance.

**Table 3.** GaitFormer direct transfer performance on gait recognition on CASIA-B and FVG datasets. We highlight in **bold** the best overall result for each dataset.

|      |             | CASIA-B      |              |              |              | FVG          |              |              |              |
|------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Size | Training    | NM           | CL           | CB           | WS           | CB           | CL           | CBG          | ALL          |
| SM   | Contrastive | 89.00        | 22.36        | 61.88        | 77.33        | 81.82        | 54.27        | 86.75        | 77.33        |
| MD   | Contrastive | 90.18        | 23.46        | 60.78        | 78.33        | 72.73        | 49.15        | 83.33        | 78.33        |
| XL   | Contrastive | 91.79        | 21.11        | 63.12        | 76.33        | 69.70        | 48.29        | 87.61        | 76.33        |
| SM   | MT          | 92.52        | 22.73        | <b>67.16</b> | 84.67        | 81.82        | <b>59.40</b> | <b>91.45</b> | 84.67        |
| MD   | MT          | <b>92.52</b> | <b>23.31</b> | 65.10        | <b>85.33</b> | <b>87.88</b> | 53.42        | 88.89        | <b>85.33</b> |
| XL   | MT          | 90.69        | 20.75        | 60.34        | 85.00        | 81.82        | 51.71        | 91.03        | 85.00        |



**Figure 6.** Fine-tuning results on gait recognition on CASIA-B and FVG, on progressively larger number of runs per identity. Compared to the same network randomly initialized, pretraining on DenseGait offers substantial improvements, even in the direct transfer regime. A consistent performance increase is obtained when also estimating attributes.

Table 4 presents state-of-the-art results compared with other skeleton-based gait recognition models. We showcase the results of GaitFormerSM trained in the Multi-Task (MT) regime, without fine-tuning (direct) and tuned with all the available training data in CASIA-B. For comparison, we include WildGait [11] with and without fine-tuning, as this model is also pretrained on a large dataset of skeleton sequences. We also compare with our implementation of GaitGraph Teepe et al. [78]—a multi-branch ST-GCN which processes joint coordinates, velocities and bone angles, achieving great results on CASIA-B—and with a ST-GCN pretrained on DenseGait.

It is clear that the fine-tuned GaitFormerSM has very good results even without fine-tuning, achieving comparable results with the state of the art. Fine-tuning marginally increases the performance, achieving 96.2% accuracy on normal walking (NM) and 72.5% performance in carry bag (CB).

**Table 4.** GaitFormer comparison to other skeleton-based gait recognition methods on CASIA-B dataset. In all methods the gallery set contains all viewpoints except the probe angle. In **bold** and underline we highlight the best and second best results for a particular viewpoint and walking condition.

| Method                 | 0°          | 18°         | 36°          | 54°         | 72°         | 90°         | 108°        | 126°         | 144°        | 162°        | 180°        | Mean        |
|------------------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|
| GaitGraph              | 79.8        | 89.5        | 91.1         | 92.7        | 87.9        | 89.5        | 94.35       | 95.1         | 92.7        | 93.5        | 80.6        | 89.7        |
| ST-GCN (DenseGait)     | 89.5        | 89.5        | 95.1         | 87.9        | 81.4        | 68.5        | 64.5        | 89.5         | 88.7        | 84.6        | 82.2        | 83.8        |
| WildGait—direct        | 72.6        | 84.6        | 90.3         | 83.8        | 63.7        | 62.9        | 66.1        | 83.0         | 86.3        | 84.6        | 83.0        | 78.3        |
| <b>NM</b> PoseFrame    | 66.9        | 90.3        | 91.1         | 55.6        | 89.5        | <b>97.6</b> | <b>98.4</b> | 97.6         | 89.5        | 69.4        | 68.5        | 83.1        |
| GaitFormerSM—MT—direct | <u>94.3</u> | <u>97.5</u> | <u>99.2</u>  | <u>98.4</u> | 79.8        | 80.6        | 89.5        | <b>100.0</b> | 94.3        | <u>95.1</u> | <u>88.7</u> | 92.5        |
| WildGait—tuned         | 86.3        | 96.0        | 97.6         | 94.3        | <b>92.7</b> | <u>94.3</u> | 94.3        | 98.4         | <u>97.6</u> | 91.1        | 83.8        | <u>93.4</u> |
| GaitFormerSM—MT—tuned  | <b>96.7</b> | <b>99.2</b> | <b>100.0</b> | <b>99.2</b> | <u>91.9</u> | 91.9        | <u>95.1</u> | <u>98.4</u>  | <b>96.7</b> | <b>97.6</b> | <b>91.1</b> | <b>96.2</b> |
| GaitGraph              | 27.4        | 33.0        | 40.3         | 37.1        | 33.8        | 33.0        | 35.4        | 33.8         | 34.6        | 21.7        | 17.7        | 31.6        |
| ST-GCN (DenseGait)     | 18.5        | 22.5        | 25.0         | 21.7        | 13.7        | 18.5        | 21.7        | 31.4         | 21.7        | 21.7        | 16.9        | 21.2        |
| WildGait—direct        | 12.1        | <u>33.0</u> | 25.8         | 18.5        | 12.9        | 11.3        | 21.7        | 24.2         | 20.1        | 26.6        | 16.1        | 20.2        |
| <b>CL</b> PoseFrame    | 13.7        | 29.0        | 20.2         | 19.4        | 28.2        | <b>53.2</b> | <b>57.3</b> | <b>52.4</b>  | 25.8        | 26.6        | 21.0        | 31.5        |
| GaitFormerSM/MT—direct | 12.9        | 21.7        | 29.0         | 25.8        | 16.1        | 18.5        | 22.5        | 29.0         | 27.4        | 26.6        | 20.1        | 22.7        |
| WildGait—tuned         | <u>29.0</u> | 32.2        | <b>35.5</b>  | <b>40.3</b> | <u>26.6</u> | 25.0        | 38.7        | 38.7         | 31.4        | <u>34.6</u> | <b>31.4</b> | <b>33.0</b> |
| GaitFormerSM/MT—tuned  | <b>35.5</b> | <b>35.5</b> | <u>33.8</u>  | <u>33.8</u> | 20.9        | 30.6        | 31.4        | 31.4         | <u>28.2</u> | <b>42.7</b> | <u>29.8</u> | <u>32.2</u> |
| GaitGraph              | 64.5        | 69.3        | 70.1         | 62.9        | 61.2        | 58.8        | 59.6        | 58.0         | 57.2        | 55.6        | 45.9        | 60.3        |
| ST-GCN (DenseGait)     | 78.2        | 68.5        | 71.7         | 60.4        | 59.6        | 45.9        | 46.7        | 58.0         | 58.0        | 58.0        | 51.6        | 59.7        |
| WildGait—direct        | 67.7        | 60.5        | 63.7         | 51.6        | 47.6        | 39.5        | 41.1        | 50.0         | 52.4        | 51.6        | 42.7        | 51.7        |
| <b>BG</b> PoseFrame    | 45.2        | 66.1        | 60.5         | 42.7        | <u>58.1</u> | <b>84.7</b> | <b>79.8</b> | <b>82.3</b>  | <u>65.3</u> | 54.0        | 50.0        | 62.6        |
| GaitFormerSM/MT—direct | <u>78.2</u> | <u>71.7</u> | <b>84.7</b>  | <b>74.2</b> | 56.4        | 50.0        | 57.2        | 66.1         | <u>69.3</u> | <u>70.9</u> | <u>59.6</u> | <u>67.1</u> |
| WildGait—fine-tuned    | 66.1        | 70.1        | 72.6         | 65.3        | 56.4        | 64.5        | 65.3        | 67.7         | 57.2        | 66.1        | 52.4        | 64.0        |
| GaitFormerSM/MT—tuned  | <b>82.2</b> | <b>80.6</b> | <u>83.8</u>  | <u>72.6</u> | <b>62.9</b> | <u>69.3</u> | <u>68.5</u> | <u>70.1</u>  | <b>69.3</b> | <b>77.4</b> | <b>60.4</b> | <b>72.5</b> |

#### 4.2. Comparison with ST-GCN and Other Pretraining Datasets

In Table 5, we compare GaitFormer with ST-GCN [26] under different pretraining datasets. Reported results are mean accuracy across all angles for CASIA-B, under normal walking (NM) scenario, and accuracy under ALL scenario for FVG. The networks were not fine-tuned on these datasets; we present direct transfer performance after pretraining. We chose to pretrain on OU-ISIR [13], as this dataset is one of the most popular, large-scale datasets for gait recognition. However, OU-ISIR lacks data diversity, as all individuals are walking on a treadmill for a short duration, which is not the case for DenseGait. We also chose to pretrain on GREW [12], as it is also a diverse dataset collected in the wild, but contains fewer identities that walk for a comparably shorter duration of time.

Results show that, as a pretraining dataset, DenseGait is consistently outperforming GREW and OU-ISIR across the two architectures. These results are consistent with the insights in Figure 1, in which we posit that longer tracking duration for the individuals imply larger data diversity when pretraining in a contrastive self-supervised fashion, which directly improves performance.

#### 4.3. Gait-Based Gender Detection

Table 6 presents results for direct transfer (zero-shot) performance for gender estimation on CASIA-gender and FVG-gender. In this case, we compared different sizes of GaitFormer trained on DenseGait in two manners: i) only estimating attributes, without a contrastive objective (Attributes Only), and ii) estimating attributes and identity using a contrastive objective (MT). Similarly to the case of gait recognition, the Multi-Task networks consistently outperforms the other training regime. Moreover, the networks achieved reasonable performance in terms of F<sub>1</sub> score (76.18% for CASIA-gender and

86.81% for FVG-gender), considering that the networks were not exposed to any manually annotated data.

**Table 5.** Comparison between GaitFormer and ST-GCN pretrained with Supervised Contrastive on GREW [12], OU-ISIR [13] and our proposed DenseGait. Performance is directly correlated with mean tracklet duration on each dataset as shown in Figure 1. We highlight in **bold** the best results for each architecture and dataset.

| Backbone                   | Pretraining Data          | CASIA-B (NM) | FVG (ALL)    |
|----------------------------|---------------------------|--------------|--------------|
| ST-GCN                     | OU-ISIR                   | 55.65        | 63.33        |
|                            | GREW                      | 61.14        | 56.67        |
|                            | DenseGait ( <b>ours</b> ) | <b>83.80</b> | <b>75.28</b> |
| GaitFormer ( <b>ours</b> ) | OU-ISIR                   | 25.73        | 51.34        |
|                            | GREW                      | 65.40        | 64.04        |
|                            | DenseGait ( <b>ours</b> ) | <b>89.0</b>  | <b>77.33</b> |

**Table 6.** GaitFormer direct transfer performance on gait-based gender estimation on CASIA-gender and FVG-gender. We highlight in **bold** the best overall results for each dataset.

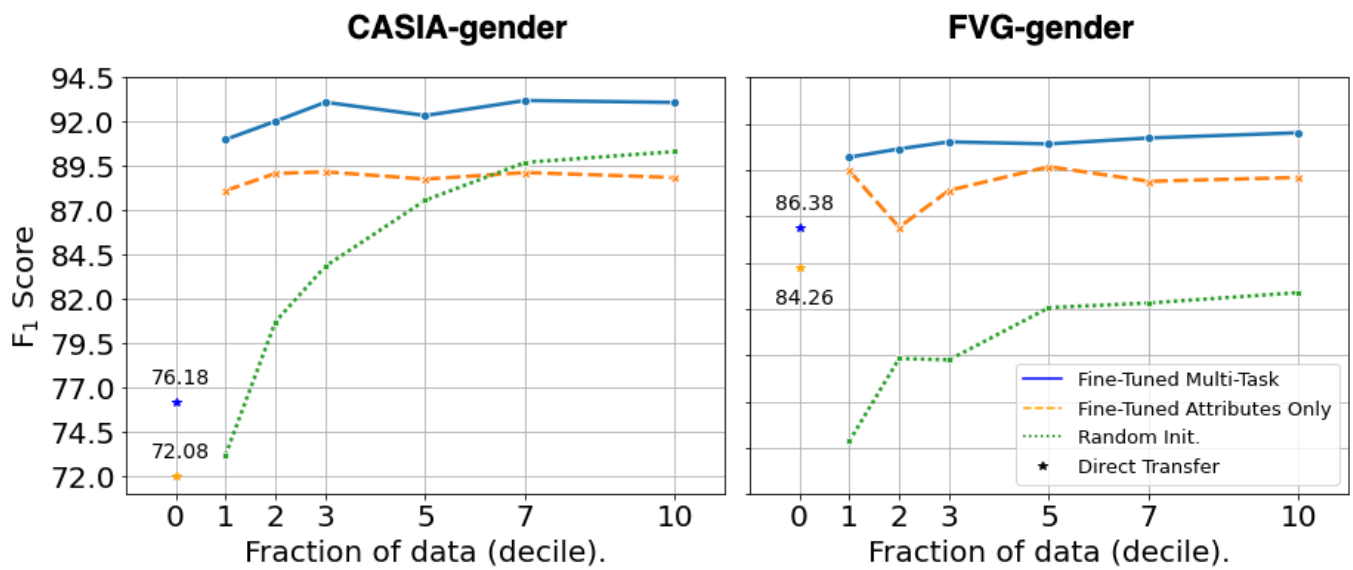
| Size | Training   | CASIA-Gender |              |                | FVG-Gender   |              |                |
|------|------------|--------------|--------------|----------------|--------------|--------------|----------------|
|      |            | Prec.        | Recall       | F <sub>1</sub> | Prec.        | Recall       | F <sub>1</sub> |
| SM   | Attributes | 94.54        | 62.46        | 72.10          | 85.87        | 84.9         | 85.10          |
| MD   | Attributes | 94.48        | 63.66        | 73.07          | 82.03        | 79.87        | 80.27          |
| XL   | Attributes | 94.59        | 62.43        | 72.08          | 84.36        | 84.43        | 84.26          |
| SM   | MT         | <b>94.84</b> | 63.61        | 73.00          | 86.47        | 86.34        | 86.33          |
| MD   | MT         | 94.71        | 61.50        | 71.31          | 86.96        | <b>86.87</b> | <b>86.81</b>   |
| XL   | MT         | 94.72        | <b>67.67</b> | <b>76.18</b>   | <b>87.06</b> | 86.21        | 86.38          |

Figure 7 presents the performance under fine-tuning of GaitFormerXL on CASIA-gender and FVG-gender, in similar conditions to the recognition task. All networks are trained with a binary-crossentropy objective on the gender estimation task, without taking the person identity into account at training time. GaitFormerXL under Multi-Task training regime is consistently superior to a network initialized from random weights, achieving an F<sub>1</sub> score of 93.09% on CASIA-gender and of 91.51% on FVG-gender. The pretrained models significantly benefit from fine-tuning when small amounts of training data is available. Performance slightly increases with the availability of more training data.

#### 4.4. Gait-Based Pedestrian Attribute Identification


For pedestrian attribute identification, we process a 10-h surveillance stream, corresponding to 10,733 tracklets, and use it for testing. For evaluation, we use the attribute pseudo-labels annotated automatically by the PAI ensemble. Figure 8 showcases R<sup>2</sup> score results for GaitFormerMD trained with a multi-task objective. This score is computed relative to the soft pseudo-labels estimated by the PAI ensemble. We emphasize that the model only uses movement information to estimate these labels, and has no information regarding appearance. Using a skeleton-based model for pedestrian attribute identification is useful in situations where the appearance of the person is unavailable (i.e., in privacy-critical scenarios). The model is effectively distilling external appearance into movement representations.

The model obtains good results in categories such as *Gender*, *AgeGroup*, *BodyType* and *Viewpoint*. The model is able to obtain better than average performance on categories such as *Footwear*, and some types of clothing. However, some clothing categories have proven to be very difficult to model, especially *LongCoat* and *Trousers*. We hypothesize that such pieces of clothing negatively affect the accuracy of the pose estimation model, resulting in low quality extracted skeletons.



**Figure 7.** Fine-tuning results for GaitFormer on CASIA-gender and FVG-gender, trained on progressively larger samples of the datasets. Compared to a randomly initialized network, GaitFormer benefits significantly from fine-tuning in extremely low data regimes (e.g., 10% of available annotated data). Compared to only pretraining on predicting attributes (Attributes Only), the Multi-Task network has consistently better performance across all fractions of the datasets.

These are promising results which show that external appearance and movement are intrinsically linked together. This is evident in the more explicit relationship between, for example, footwear and gait, in which, intuitively, gait is severely affected by the walker’s choice of shoes. Clothing, accessories, and actions while walking can be regarded as “distractor” attributes, which affect gait only temporarily. However, there are more subtle information cues which are present in gait, related to the developmental aspects of the person (e.g., gender, age, body composition, mental state etc). These attributes are more stable in time, and can provide insights into the internal workings of the walker. We posit that, in the future, works in gait analysis will tackle more rigorously the problem of estimating the internal state of the walker (i.e., personality/mental issues) through specialized datasets and methods.

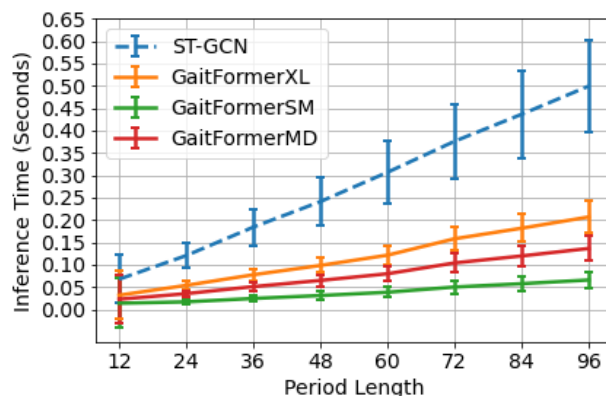


**Figure 8.** GaitFormerMD performance in terms of  $R^2$  score. GaitFormerMD was trained with the multi-task objective. The model uses only movement information to predict attributes, and no information regarding the appearance of the individual.

#### 4.5. Inference Time

Using transformer architectures for processing gait has other advantages besides a noticeable increase in downstream performance. Transformers have been shown to be more efficient in terms of inference time when compared to convolutional networks [56]. This effect is not directly correlated with the number of parameters, but is rather more influenced by the network structure [79].

In Figure 9, we show a comparison between multiple sizes of GaitFormer, a plain transformer module minimally adapted for processing skeleton sequences, with the ST-GCN network, a popular architecture for skeleton action recognition [26] and gait analysis [46]. We computed the inference time across multiple period lengths (from 12 frames to 96 frames) to evaluate the scalability when processing shorter/longer sequences. For each period length, we run 100 experiments with a batch size of 512 and show the mean inference time in seconds, along with the standard deviation. All experiments were run on a NVIDIA RTX 3060 GPU. Even with comparable and exceeding number of parameters (ST-GCN from Cosma and Radoi [11] has 3.11M parameters), the transformer architecture clearly outperforms graph-convolutional models for processing gait sequences across multiple sequence lengths.



**Figure 9.** Inference times across processed walking duration length (period length) for ST-GCN and the various sizes of GaitFormer. We report the mean and standard deviation across 100 runs, for each period length.

## 5. Conclusions

In this work, we presented DenseGait, currently the largest dataset for pretraining gait analysis models, consisting of 217 K anonymized skeleton sequences. Each skeleton sequence is automatically annotated with 42 appearance attributes by making use of an ensemble of pretrained PAI networks. We make DenseGait available to the research community, under open credentialized access, to promote further advancement in the skeleton-based gait analysis field. We proposed GaitFormer, a transformer that is pre-trained on DenseGait in a self-supervised and multi-task fashion. The model obtains 92.5% accuracy on CASIA-B and 85.3% accuracy on FVG, without processing any manually annotated data, achieving higher performance even compared to fully supervised methods. GaitFormer represents the first application of plain transformer encoders for skeleton-based gait analysis, without any hand-crafted architectural modifications. We explored pedestrian attribute identification based solely on movement, without utilizing appearance information. GaitFormer achieves good results in gender, age body type, and clothing attributes.

**Author Contributions:** Conceptualization, A.C and E.R.; Formal analysis, A.C.; Funding acquisition, E.R.; Methodology, A.C. and E.R.; Project administration, E.R.; Resources, E.R.; Software, A.C.; Supervision, E.R.; Validation, A.C.; Writing—original draft, A.C.; Writing—review & editing, E.R. All authors have read and agreed to the published version of the manuscript.



**Funding:** This work was partly supported by CRC Research Grant 2021, with funds from UEFISCDI in project CORNET (PN-III 1/2018) and by the Google IoT/Wearables Student Grants.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The dataset analyzed in this study are made publicly available. This data can be found here: <https://bit.ly/3SLO8RW>, accessed on 1 July 2022.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

- Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; Hoi, S.C. Deep Learning for Person Re-identification: A Survey and Outlook. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *44*, 2872–2893. [CrossRef] [PubMed]
- Wang, M.; Deng, W. Deep face recognition: A survey. *Neurocomputing* **2021**, *429*, 215–244. [CrossRef]
- Nixon, M.S.; Tan, T.N.; Chellappa, R. *Human Identification Based on Gait (The Kluwer International Series on Biometrics)*; Springer: Berlin/Heidelberg, Germany, 2005.
- McGibbon, C.A. Toward a Better Understanding of Gait Changes With Age and Disablement: Neuromuscular Adaptation. *Exerc. Sport Sci. Rev.* **2003**, *31*, 102–108. [CrossRef] [PubMed]
- Kumar, R.; Isik, C.; Phoha, V.V. Treadmill Assisted Gait Spoofing (TAGS) An Emerging Threat to Wearable Sensor-based Gait Authentication. *Digit. Threat. Res. Pract.* **2021**, *2*, 1–17. [CrossRef]
- Singh, J.P.; Jain, S.; Arora, S.; Singh, U.P. Vision-based gait recognition: A survey. *IEEE Access* **2018**, *6*, 70497–70527. [CrossRef]
- Makihara, Y.; Nixon, M.S.; Yagi, Y. Gait recognition: Databases, representations, and applications. In *Computer Vision: A Reference Guide*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 1–13.
- Zhang, Z.; Tran, L.; Liu, F.; Liu, X. On Learning Disentangled Representations for Gait Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *44*, 345–360. [CrossRef]
- Sepas-Moghaddam, A.; Etemad, A. View-invariant gait recognition with attentive recurrent learning of partial representations. *IEEE Trans. Biom. Behav. Identity Sci.* **2020**, *3*, 124–137. [CrossRef]
- Thapar, D.; Nigam, A.; Aggarwal, D.; Agarwal, P. VGR-net: A view invariant gait recognition network. In Proceedings of the 2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA), Singapore, 11–12 January 2018; pp. 1–8.
- Cosma, A.; Radoi, I.E. WildGait: Learning Gait Representations from Raw Surveillance Streams. *Sensors* **2021**, *21*, 8387. [CrossRef]
- Zhu, Z.; Guo, X.; Yang, T.; Huang, J.; Deng, J.; Huang, G.; Du, D.; Lu, J.; Zhou, J. Gait Recognition in the Wild: A Benchmark. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Montreal, BC, Canada, 11–17 October 2021.
- Makihara, Y.; Mannami, H.; Tsuji, A.; Hossain, M.; Sugiura, K.; Mori, A.; Yagi, Y. The OU-ISIR Gait Database Comprising the Treadmill Dataset. *IPSJ Trans. Comput. Vis. Appl.* **2012**, *4*, 53–62. [CrossRef]
- Xu, C.; Makihara, Y.; Ogi, G.; Li, X.; Yagi, Y.; Lu, J. The OU-ISIR Gait Database Comprising the Large Population Dataset with Age and Performance Evaluation of Age Estimation. *IPSJ Trans. Comput. Vis. Appl.* **2017**, *9*, 1–14. [CrossRef]
- Yu, S.; Tan, D.; Tan, T. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In Proceedings of the 18th International Conference on Pattern Recognition (ICPR'06), Hong Kong, China, 20–24 August 2006; Volume 4, pp. 441–444.
- Liu, X.; Zhao, H.; Tian, M.; Sheng, L.; Shao, J.; Yi, S.; Yan, J.; Wang, X. Hydraplus-net: Attentive deep features for pedestrian analysis. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 350–359.
- Tang, C.; Sheng, L.; Zhang, Z.; Hu, X. Improving pedestrian attribute recognition with weakly-supervised multi-scale attribute-specific localization. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Korea, 27–28 October 2019; pp. 4997–5006.
- Jian, J.; Houjing, H.; Wenjie, Y.; Xiaotang, C.; Kaiqi, H. Rethinking of pedestrian attribute recognition: Realistic datasets with efficient method. *arXiv* **2020**, arXiv:2005.11909.
- Catrana, A.; Cosma, A.; Radoi, I.E. From Face to Gait: Weakly-Supervised Learning of Gender Information from Walking Patterns. In Proceedings of the 2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021), Jodhpur, India, 15–18 December 2021; pp. 1–5.
- uk Ko, S.; Tolea, M.I.; Hausdorff, J.M.; Ferrucci, L. Sex-specific differences in gait patterns of healthy older adults: Results from the Baltimore Longitudinal Study of Aging. *J. Biomech.* **2011**, *44*, 1974–1979. [CrossRef]
- Ko, S.u.; Hausdorff, J.M.; Ferrucci, L. Age-associated differences in the gait pattern changes of older adults during fast-speed and fatigue conditions: results from the Baltimore longitudinal study of ageing. *Age Ageing* **2010**, *39*, 688–694. [CrossRef]
- Choi, H.; Lim, J.; Lee, S. Body fat-related differences in gait parameters and physical fitness level in weight-matched male adults. *Clin. Biomech.* **2021**, *81*, 105243. [CrossRef]
- Takayanagi, N.; Sudo, M.; Yamashiro, Y.; Lee, S.; Kobayashi, Y.; Niki, Y.; Shimada, H. Relationship between daily and in-laboratory gait speed among healthy community-dwelling older adults. *Sci. Rep.* **2019**, *9*, 3496. [CrossRef]
- Liu, Z.; Malave, L.; Osuntogun, A.; Sudhakar, P.; Sarkar, S. Toward understanding the limits of gait recognition. In Proceedings of the Biometric Technology for Human Identification, Orlando, FL, USA, 12–13 April 2004; Volume 5404, pp. 195–205.

25. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is All you Need. In *Proceedings of the Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30.
26. Yan, S.; Xiong, Y.; Lin, D. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, New Orleans, LA, USA, 2–7 February 2018.
27. Plizzari, C.; Cannici, M.; Matteucci, M. Spatial temporal transformer network for skeleton-based action recognition. In *Proceedings of the International Conference on Pattern Recognition*, Virtual Event, 10–15 January 2021; pp. 694–701.
28. Hofmann, M.; Geiger, J.; Bachmann, S.; Schuller, B.; Rigoll, G. The tum gait from audio, image and depth (gaid) database: Multimodal recognition of subjects and traits. *J. Vis. Commun. Image Represent.* **2014**, *25*, 195–206. [[CrossRef](#)]
29. Chen, T.; Kornblith, S.; Norouzi, M.; Hinton, G. A simple framework for contrastive learning of visual representations. In *Proceedings of the International Conference on Machine Learning*, PMLR, Virtual Event, 13–18 July 2020; pp. 1597–1607.
30. Shutler, J.D.; Grant, M.G.; Nixon, M.S.; Carter, J.N. On a large sequence-based human gait database. In *Applications and Science in Soft Computing*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 339–346.
31. Sarkar, S.; Phillips, P.J.; Liu, Z.; Vega, I.R.; Grother, P.; Bowyer, K.W. The humanid gait challenge problem: Data sets, performance, and analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **2005**, *27*, 162–177. [[CrossRef](#)]
32. Do, T.D.; Nguyen, V.H.; Kim, H. Real-time and robust multiple-view gender classification using gait features in video surveillance. *Pattern Anal. Appl.* **2020**, *23*, 399–413. [[CrossRef](#)]
33. Xu, S.; Fang, J.; Hu, X.; Ngai, E.; Guo, Y.; Leung, V.; Cheng, J.; Hu, B. Emotion recognition from gait analyses: Current research and future directions. *arXiv* **2020**, arXiv:2003.11461.
34. Zhang, D.; Wang, Y.; Bhanu, B. Ethnicity classification based on gait using multi-view fusion. In *Proceedings of the 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*, San Francisco, CA, USA, 13–18 June 2010; pp. 108–115.
35. Sepas-Moghaddam, A.; Etemad, A. Deep gait recognition: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**. [[CrossRef](#)]
36. Han, J.; Bhanu, B. Individual Recognition Using Gait Energy Image. *IEEE Trans. Pattern Anal. Mach. Intell.* **2006**, *28*, 316–322. [[CrossRef](#)]
37. Wang, C.; Zhang, J.; Pu, J.; Yuan, X.; Wang, L. Chrono-Gait Image: A Novel Temporal Template for Gait Recognition. In *Proceedings of the Computer Vision—ECCV 2010*; Daniilidis, K.; Maragos, P., Paragios, N., Eds.; Springer: Berlin/Heidelberg, Germany, 2010; pp. 257–270.
38. Bashir, K.; Xiang, T.; Gong, S. Gait recognition using Gait Entropy Image. In *Proceedings of the 3rd International Conference on Imaging for Crime Detection and Prevention (ICDP 2009)*, London, UK, 3 December 2009; pp. 1–6. [[CrossRef](#)]
39. Chao, H.; He, Y.; Zhang, J.; Feng, J. Gaitset: Regarding gait as a set for cross-view gait recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 8126–8133.
40. Fan, C.; Peng, Y.; Cao, C.; Liu, X.; Hou, S.; Chi, J.; Huang, Y.; Li, Q.; He, Z. Gaitpart: Temporal part-based model for gait recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Seattle, WA, USA, 14–19 June 2020; pp. 14225–14233.
41. Lin, B.; Zhang, S.; Yu, X. Gait recognition via effective global-local feature representation and local temporal aggregation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, Montreal, QC, Canada, 11–17 October 2021; pp. 14648–14656.
42. Li, J.; Wang, C.; Zhu, H.; Mao, Y.; Fang, H.S.; Lu, C. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, USA, 15–20 June 2019; pp. 10863–10872.
43. Cosma, A.; Radoi, I.E. Multi-Task Learning of Confounding Factors in Pose-Based Gait Recognition. In *Proceedings of the 2020 19th RoEduNet Conference: Networking in Education and Research (RoEduNet)*, Bucharest, Romania, 11–12 December 2020; pp. 1–6. [[CrossRef](#)]
44. Liao, R.; Cao, C.; Garcia, E.B.; Yu, S.; Huang, Y. Pose-Based Temporal-Spatial Network (PTSN) for Gait Recognition with Carrying and Clothing Variations. In *Proceedings of the Biometric Recognition*, Shenzhen, China, 28–29 October 2017; pp. 474–483.
45. An, W.; Liao, R.; Yu, S.; Huang, Y.; Yuen, P.C. Improving Gait Recognition with 3D Pose Estimation. In *Proceedings of the CCB, Urumqi, China*, 11–12 August 2018.
46. Li, N.; Zhao, X.; Ma, C. JointsGait: A model-based Gait Recognition Method based on Gait Graph Convolutional Networks and Joints Relationship Pyramid Mapping. *arXiv* **2020**, arXiv:2005.08625.
47. Li, N.; Zhao, X. A Strong and Robust Skeleton-based Gait Recognition Method with Gait Periodicity Priors. *IEEE Trans. Multimed.* **2022**. [[CrossRef](#)]
48. Xu, C.; Makihara, Y.; Li, X.; Yagi, Y.; Lu, J. Gait recognition from a single image using a phase-aware gait cycle reconstruction network. In *Proceedings of the European Conference on Computer Vision*, Virtual Event, 23–28 August 2020; pp. 386–403.
49. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*; Association for Computational Linguistics: Minneapolis, MN, USA, 2019; pp. 4171–4186. [[CrossRef](#)]

50. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. In *Proceedings of the Advances in Neural Information Processing Systems*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2020; Volume 33, pp. 1877–1901.
51. Gidaris, S.; Singh, P.; Komodakis, N. Unsupervised Representation Learning by Predicting Image Rotations. In *Proceedings of the International Conference on Learning Representations*, Vancouver, Canada, 30 April–3 May 2018.
52. Doersch, C.; Gupta, A.; Efros, A.A. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, Santiago, Chile, 7–13 December 2015; pp. 1422–1430.
53. Wang, J.; Jiao, J.; Liu, Y.H. Self-supervised video representation learning by pace prediction. In *Proceedings of the European Conference on Computer Vision*, Virtual Event, 23–28 August 2020; pp. 504–521.
54. Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; Deny, S. Barlow twins: Self-supervised learning via redundancy reduction. In *Proceedings of the International Conference on Machine Learning*, PMLR, Virtual Event, 18–24 July 2021; pp. 12310–12320.
55. Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; Joulin, A. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, Virtual Event, 11–17 October 2021; pp. 9650–9660.
56. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An Image is Worth 16 × 16 Words: Transformers for Image Recognition at Scale. In *Proceedings of the International Conference on Learning Representations*, Virtual Event, 26–30 April 2020.
57. Zhang, H.; Hao, Y.; Ngo, C.W. Token shift transformer for video classification. In *Proceedings of the 29th ACM International Conference on Multimedia*, Chengdu, China, 20–24 October 2021; pp. 917–925.
58. Dong, L.; Xu, S.; Xu, B. Speech-Transformer: A No-Recurrence Sequence-to-Sequence Model for Speech Recognition. In *Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, 15–20 April 2018; pp. 5884–5888. [[CrossRef](#)]
59. Beal, J.; Wu, H.Y.; Park, D.H.; Zhai, A.; Kislyuk, D. Billion-Scale Pretraining with Vision Transformers for Multi-Task Visual Representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, Waikoloa, HI, USA, 4–8 January 2022; pp. 564–573.
60. Lin, T.Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; Zitnick, C.L. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, Zurich, Switzerland, 6–12 September 2014; pp. 740–755.
61. Nasser, M.H.; Moradi, H.; Hosseini, R.; Babaei, M. Simple online and real-time tracking with occlusion handling. *arXiv* **2021**, arXiv:2103.04147.
62. Bewley, A.; Ge, Z.; Ott, L.; Ramos, F.; Upcroft, B. Simple online and realtime tracking. In *Proceedings of the 2016 IEEE International Conference on Image Processing (ICIP)*, Phoenix, AZ, USA, 25–28 September 2016; pp. 3464–3468.
63. Wojke, N.; Bewley, A.; Paulus, D. Simple Online and Realtime Tracking with a Deep Association Metric. In *Proceedings of the 2017 IEEE International Conference on Image Processing (ICIP)*, Beijing, China, 17–20 September 2017; pp. 3645–3649. [[CrossRef](#)]
64. Murray, M.P.; Drought, A.B.; Kory, R.C. Walking Patterns of Normal Men. *JBJS* **1964**, *46*, 335–360. [[CrossRef](#)]
65. Xu, X.; Meng, Q.; Qin, Y.; Guo, J.; Zhao, C.; Zhou, F.; Lei, Z. Searching for alignment in face recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Virtual Event, 19 August 2021; Volume 35, pp. 3065–3073.
66. Li, X.; Makihara, Y.; Xu, C.; Yagi, Y.; Ren, M. Gait-based human age estimation using age group-dependent manifold learning and regression. *Multimed. Tools Appl.* **2018**, *77*, 28333–28354. [[CrossRef](#)]
67. Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, USA, 26 June–1 July 2016; pp. 2818–2826.
68. Li, D.; Zhang, Z.; Chen, X.; Huang, K. A richly annotated pedestrian dataset for person retrieval in real surveillance scenarios. *IEEE Trans. Image Process.* **2019**, *28*, 1575–1590. [[CrossRef](#)] [[PubMed](#)]
69. Deng, Y.; Luo, P.; Loy, C.C.; Tang, X. Pedestrian attribute recognition at far distance. In *Proceedings of the 22nd ACM International Conference on Multimedia*, Orlando, FL, USA, 3–7 November 2014; pp. 789–792.
70. Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C.A.; Cubuk, E.D.; Kurakin, A.; Li, C.L. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 596–608.
71. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)]
72. Gabeur, V.; Sun, C.; Alahari, K.; Schmid, C. Multi-modal transformer for video retrieval. In *Proceedings of the European Conference on Computer Vision*, Virtual Event, 23–28 August 2020; pp. 214–229.
73. Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; Krishnan, D. Supervised Contrastive Learning. In *Proceedings of the Advances in Neural Information Processing Systems*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2020; Volume 33, pp. 18661–18673.
74. Müller, R.; Kornblith, S.; Hinton, G.E. When does label smoothing help? *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 4694–4703.
75. Hinton, G.; Vinyals, O.; Dean, J. Distilling the Knowledge in a Neural Network. In *Proceedings of the NIPS Deep Learning and Representation Learning Workshop*, Montreal, QC, Canada, 11–12 December 2015.
76. Smith, L.N. Cyclical learning rates for training neural networks. In *Proceedings of the 2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Santa Rosa, CA, USA, 24–31 March 2017; pp. 464–472.
77. Zhang, T.; Wu, F.; Katiyar, A.; Weinberger, K.Q.; Artzi, Y. Revisiting Few-sample BERT Fine-tuning. In *Proceedings of the International Conference on Learning Representations*, Virtual Event, 26–30 April 2020.

- 
78. Teepe, T.; Khan, A.; Gilg, J.; Herzog, F.; Hörmann, S.; Rigoll, G. GaitGraph: Graph Convolutional Network for Skeleton-Based Gait Recognition. In Proceedings of the 2021 IEEE International Conference on Image Processing (ICIP), Anchorage, AK, USA, 19–22 September 2021; pp. 2314–2318. [[CrossRef](#)]
  79. Langerman, D.; Johnson, A.; Buettner, K.; George, A.D. Beyond Floating-Point Ops: CNN Performance Prediction with Critical Datapath Length. In Proceedings of the 2020 IEEE High Performance Extreme Computing Conference (HPEC), Waltham, MA, USA, 22–24 September 2020; pp. 1–9. [[CrossRef](#)]