

On Model and Data Scaling for Skeleton-based Self-Supervised Gait Recognition

Adrian Cosma^{1,2}, Andy Eduard Cătrună¹, Emilian Rădoi¹

¹National University of Science and Technology POLITEHNICA Bucharest

²Dalle Molle Institute for Artificial Intelligence Research (IDSIA)

{ioan.adrian.cosma, andy_eduard.catruna, emilian.radoi}@upb.ro

Abstract

Gait recognition from video streams is a challenging problem in computer vision biometrics due to the subtle differences between gaits and numerous confounding factors. Recent advancements in self-supervised pretraining have led to the development of robust gait recognition models that are invariant to walking covariates. While neural scaling laws have transformed model development in other domains by linking performance to data, model size, and compute, their applicability to gait remains unexplored. In this work, we conduct the first empirical study scaling on skeleton-based self-supervised gait recognition to quantify the effect of data quantity, model size and compute on downstream gait recognition performance. We pretrain multiple variants of GaitPT – a transformer-based architecture – on a dataset of 2.7 million walking sequences collected in the wild. We evaluate zero-shot performance across four benchmark datasets to derive scaling laws for data, model size, and compute. Our findings demonstrate predictable power-law improvements in performance with increased scale and confirm that data and compute scaling significantly influence downstream accuracy. We further isolate architectural contributions by comparing GaitPT with GaitFormer under controlled compute budgets. These results provide practical insights into resource allocation and performance estimation for real-world gait recognition systems.

Introduction

Modern AI systems scale predictably: more data, more parameters, better performance (Kaplan et al. 2020; Hoffmann et al. 2022; Caballero et al. 2023; Hernandez et al. 2021; Muennighoff et al. 2023; Yao and Wang 2023). But does this hold for gait — one of the most subtle and privacy-sensitive biometric modalities? Gait recognition from video streams is a long-standing and difficult problem in the field of computer vision biometrics, due to the subtle differences between gaits across individuals, as well as the innumerable amounts of confounding factors in processing walks (Nixon, Tan, and Chellappa 2005). A person’s gait is affected, for example, by their clothing, accessories, psychological state and action performed during walking. Furthermore, external factors such as the weather and data acquisition hardware introduce additional measurement errors, as

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

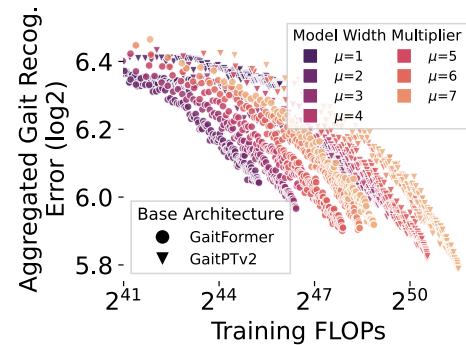


Figure 1: We trained multiple scales of skeleton-based gait recognition models using self-supervised learning on 2.7M walking sequences and analysed scaling trends in terms of model size, dataset size and compute allocation.

scene lighting, the subject’s distance from the camera, video framerate, and viewpoint affect the final representation of the walking sequence. A large body of work has been devoted to explicitly isolate confounding factors, often by the means of constructing specialized architectures (Fu et al. 2023; Teepe et al. 2021; Cosma and Radoi 2022; Zhang et al. 2019) or by constructing dedicated training datasets (Cosma and Radoi 2024; Yu, Tan, and Tan 2006; Zou et al. 2024; Makihara et al. 2012; Zhang et al. 2019). However, for practical deployment of robust gait recognition models, general methods of self-supervised pretraining on an automatically constructed dataset have been developed (Cosma and Radoi 2021, 2022), which are invariant to walking covariates by sheer exposure to many different walking registers.

More recently, across domains, there has been extensive interest of methods to improve results without the construction of specialized neural architectures, but only through scaling up data and computational resources (Cherti et al. 2023; Kaplan et al. 2020; Rosenfeld et al. 2019; Cherti et al. 2023; Caballero et al. 2022). Consequently, ”the bitter lesson” in statistical representation learning (Sutton 2019) is that such purely compute-intensive approaches have been vastly outperforming other methods that explicitly integrate human knowledge. These attempts have led to the empirical discovery and mathematical formulation of neural scaling

laws (Kaplan et al. 2020), in which the relationship between model performance and amount of data or compute is expressed as a power law of the form $L \propto N^\alpha$, where L is the model loss or a performance metric such as accuracy, N is the number of parameters or amount of data and α is constant parameters found through curve fitting. Scaling laws are crucial from a practical standpoint as they enable estimation of model performance as a function of compute and data before actual training. This approach leads to reduced costs and better resource allocation, justifying whether potentially costly business decisions of scaling up data or compute are worth the performance improvement.

To date, a large scale study on the scaling behaviour of self-supervised gait recognition models has not been performed. In this work, we conduct the first comprehensive scaling study of skeleton-based self-supervised gait recognition, quantifying the effects of model size, data volume, and compute allocation on downstream zero-shot performance. For our study, we collected a dataset of 2.7M skeleton sequences from real-world video streams — the largest such dataset to date — which offers empirical scaling trends grounded in real-world variability. We study scaling properties of skeleton-based gait recognition, as human poses are lightweight, invariant to clothing or walking variations and encode mostly movement patterns as opposed to appearance information (Cătrună, Cosma, and Rădoi 2024). We benchmark a modified version of the state-of-the-art skeleton-based gait processing transformer architecture, GaitPT (Catruna, Cosma, and Radoi 2023), to utilize the latest improvements in transformer models in terms of scalability and training stability. Furthermore, we directly compare scaling performance against GaitFormer (Cosma and Radoi 2022), a strong transformer model, to isolate architectural contributions from raw scale. We analyze compute efficiency by training all models under controlled FLOP budgets and report iso-compute comparisons. We evaluate zero-shot performance on aggregated metrics from 4 different datasets: CASIA-B (Yu, Tan, and Tan 2006), PsyMo (Cosma and Radoi 2024), GREW (Zhu et al. 2021) and Gait3D (Zheng et al. 2022), to have a comprehensive view of scaling properties across controlled and in-the-wild gait recognition scenarios.

Our work makes the following contributions:

1. We provide **data scaling laws** for self-supervised gait recognition by pretraining on a dataset of 2.7M gait sequences which we collected from in-the-wild video streams. We show that downstream model performance on controlled and in-the-wild gait recognition benchmarks can be extrapolated from small scale data-bound experiments, supporting the use of power-law scaling in this domain. We also show that quality improvements have a measurable positive impact on downstream performance.
2. We derive **model scaling laws** for self-supervised gait recognition, showing predictable improvements with increasing model size. Our experiments with GaitPTv2 span multiple model scales and dataset sizes. As far as we know, we have the first reproduction of zero-shot hy-

perparameter transfer, through μP (Yang et al. 2022), in area of self-supervised gait recognition.

3. We perform a detailed **compute analysis** of model training, comparing different scaling regimes across FLOP budgets. We show that GaitPTv2 outperforms GaitFormer (Cosma and Radoi 2022) due to its increased use of compute for the same number of parameters.

Related Work

Scaling laws have been observed in statistical learning since Cortes et al. (Cortes et al. 1993) proposed power laws for model performance as a function of data size. However, scaling behaviour has only recently been extensively studied, with the availability of compute and internet-scale datasets. Notably, scaling laws have been extensively explored in large language model training (Kaplan et al. 2020; Hoffmann et al. 2022; Caballero et al. 2023; Hernandez et al. 2021; Muennighoff et al. 2023; Yao and Wang 2023), establishing a power law relationship between the number of tokens in the dataset and number of model parameters. Hoffmann et al., (Hoffmann et al. 2022) formulated scaling laws for compute-optimal training, identifying over-training and under-training regimes for language models, given a compute budget. Furthermore, Hernandez et al. (Hernandez et al. 2021) formulated scaling laws for transfer learning, in which there is a predictable downstream performance in terms of the ratio between the amount of pretraining data and amount of fine-tuning data. In general, scaling laws fall under two broad categories of *data scaling laws* and *model scaling laws*.

Aside from language modelling, scaling laws have been established for other domains as well, for example, for machine translation (Bansal et al. 2022), masked image modelling (Xie et al. 2023), contrastive language-image learning (Cherti et al. 2023) and recommendation models (Shin et al. 2023). For computer vision (Oquab et al. 2023), the need for scaling informed further architectural developments in vision transformers (Dehghani et al. 2023; Alabdulmohsin et al. 2024) for efficient distributed training.

Regarding gait processing, scaling analysis has not been extensively performed to date. Several large-scale datasets have been proposed (Zou et al. 2024; Zhu et al. 2021; Cosma and Radoi 2022; Fan et al. 2023), but lack the magnitude and diversity for properly studying data scaling in realistic environments. Cosma et al., (Cosma and Radoi 2023) combined multiple existing gait datasets into a larger set of 800k sequences for training an autoencoder-type model. However, they did not further explore the effect of the data scale. Previously, works in self-supervised learning have noticed improved downstream performance with increase in data scale (Cosma and Radoi 2021, 2022), but the amount of data is relatively small to provide insights into scaling behaviour. GaitLU-1M (Fan et al. 2023) is a large-scale in-the-wild dataset for gait recognition, but its focus is on modeling gait using sequences of silhouettes, while we are interested in analyzing scaling trends for skeleton-based gait recognition.

In this work, we provide an empirical analysis of scaling behaviour for both data size, parameter count and com-

Dataset	# IDs	# Seq.	# Covariates	# Views	Type	Env.	Annotation
FVG (Zhang et al. 2019)	226	2,857	5	3	Controlled	outdoor	laboratory [†]
CASIA-B (Yu, Tan, and Tan 2006)	124	13,640	3	11	Controlled	indoor	laboratory [†]
PsyMo (Cosma and Radoi 2024)	312	14,976	7	6	Controlled	indoor	laboratory [†]
OU-ISIR (Makihara et al. 2012)	10,307	144,298	1	14	Controlled	indoor	laboratory [†]
CCGR (Zou et al. 2024)	970	1,580,617	53	33	Controlled	indoor	laboratory [†]
Gait3D (Zheng et al. 2022)	4,000	25,309	–	–	In the Wild	indoor	manually labeled
GREW (Zhu et al. 2021)	26,000	128,000	–	–	In the Wild	outdoor	manually labeled
UWG (Cosma and Radoi 2021)	38,502	38,502	–	–	In the Wild	outdoor	unlabeled
DenseGait (Cosma and Radoi 2022)	217,954	217,954	–	–	In the Wild	outdoor	unlabeled
GaitLU-1M (Fan et al. 2023)	1,035,309	1,035,309	–	–	In the Wild	outdoor	unlabeled
Ours	2,779,774	2,779,774	–	–	In the Wild	outdoor	unlabeled

Table 1: Comparison of popular datasets for gait recognition. Our dataset, GREW, Gait3D, and DenseGait are collected in the wild while datasets marked with “†” are constructed in controlled laboratory environments.

pute utilization for gait recognition in the regime of self-supervised contrastive learning, on a large-scale in-the-wild dataset of 2.7M walking sequences.

Experimental Setup

In this section we detail our self-supervised pretraining setup, in which we describe the collection of a large unlabeled gait recognition dataset, training configuration, the choice of transformer architectures used for gait processing, model evaluation and modeling power-law relationship between downstream performance and scaling parameters.

Collection of the Pretraining Dataset

Since one of our goals is to estimate the effect of scaling up unlabeled gait data to the zero-shot performance of gait recognition models across different model scales, we gather a substantial gait recognition dataset used for self-supervised pretraining. Manually annotating gait sequences requires extensive labor, and is unfeasible for scaling up data to multiple orders of magnitude. We aim to learn informative walking representations without manual annotations.

In a manner similar to other approaches (Cosma and Radoi 2021, 2022), we process publicly available outdoor video streams, each containing a considerable amount of people walking. The videos are chosen to have a diverse range of weather conditions, times of day, camera viewpoints, geographic locations, and containing both static and moving cameras. Works in gait recognition employ either sequences of silhouettes (Chao et al. 2021; Lin et al. 2022), body meshes (Li et al. 2020; Zheng et al. 2022) or sequences of skeletons to encode walking (Teepe et al. 2021; Catruna, Cosma, and Radoi 2023; Cosma and Radoi 2022). We chose to use skeletons as they are easy to accurately extract (Fang et al. 2022; Xu et al. 2022), lightweight in terms of storage and processing, and methods using skeleton sequences have shown promising results in this area (Cosma and Radoi 2022; Catruna, Cosma, and Radoi 2023; Fu et al. 2023; Fan et al. 2024). Skeletons enable fine-grained control on data quality by offering information for each joint across time. While there are methods that also incorporate skeleton maps (Fan et al. 2024) and SMPL body meshes (Zheng et al. 2022), we opted for the simplest case in which only sequences of skeletons are used to compute gait representations.

We process the stream to extract human poses using AlphaPose (Fang et al. 2022) and track each pose in the video using SortOH (Nasseri et al. 2021). We employed minimal filtering of the extracted skeleton sequences, by keeping only sequences of a minimum of 48 frames (at a framerate of 24fps), above an average joint confidence threshold of 0.6.

Our dataset contains a diverse range of walking registers; for instance, pedestrians are walking wearing different pieces of clothing or footwear, walking while carrying luggage or shopping bags, walking alongside other people, walking while talking on the phone or doing other actions. Each person’s identity is anonymized: we discard any appearance cues and metadata and keep only the skeleton sequence. In total, our dataset contains 2.7M skeleton sequences, with an average walking duration of 168 frames for a total of 132,931 days of walking. It is currently the largest in-the-wild and unlabeled dataset reported in literature, having an order of magnitude more skeleton sequences compared to previous in-the-wild datasets (Zhu et al. 2021; Cosma and Radoi 2022). Table 1 shows a comparison with other gait recognition datasets from literature. Our dataset is private and we only use it for unsupervised pretraining to gauge the effect of scaling data on downstream zero-shot performance.

Contrastive Self-Supervised Gait Recognition

For self-supervision, a natural pretraining regime for this domain is contrastive learning, in which the model learns to separate walking sequences of different people and aggregate walks of the same person. This approach to gait recognition has been done in the past for both label and unlabeled datasets (Catruna, Cosma, and Radoi 2023; Cosma and Radoi 2022). In particular, we adopt the SimCLR (Chen et al. 2020) approach for contrastive learning, in which we augment a walking sequence in two different ways to form positive pairs. Following previous works (Cosma and Radoi 2022; Catruna, Cosma, and Radoi 2023), we used the following augmentations for skeleton sequences: random temporal crops, random flips, random mirror, joint noise, random paces and randomly smoothing the sequence. We translate and scale each sequence based on the skeleton in the middle of the sequence, adopting the “sequence normalization” approach formulated by Catruna et al. (Cătrună, Cosma, and Rădoi 2024), since it has minimal impact upon in-the-wild gait recognition scenarios. Compared to other

works (Fu et al. 2023), we do not use any explicit anthropomorphic features (e.g., limb lengths, limb angles) that may provide shortcuts and artificially increase recognition performance (Cătrună, Cosma, and Rădoi 2024).

Additionally, alongside the contrastive loss, we use the KoLeo regularizer (Sablayrolles et al. 2019) to induce uniform feature spreading within a batch. The KoLeo regularizer improves performance in retrieval-type tasks (Sablayrolles et al. 2019), and has been used in self-supervised learning on images (Oquab et al. 2023; Caron et al. 2020). Given a batch of k feature vectors (f_1, f_2, \dots, f_k) , it is defined as $\mathcal{L}_{koleo} = -\frac{1}{n} \sum_{i=1}^n \log(d_{n,i})$, where $d_{n,i} = \min_{j \neq i} \|x_j - x_i\|$, the minimum distance between x_i and each of the other vectors within the batch. As such, the pretraining loss in our setting is defined as $\mathcal{L} = \mathcal{L}_{\text{SimCLR}} + \lambda \mathcal{L}_{koleo}$, where we chose $\lambda = 0.01$.

Transformer Architectures for Gait Processing

	Model Name	Depth	d_{model}	Output	# Params	GFLOPs
Deep & Thin	GaitPTv2-1	{2, 2, 12, 2}	{4, 8, 16, 32}	32	154,696	0.036
	GaitPTv2-2		{8, 16, 32, 64}	64	607,484	0.145
	GaitPTv2-3		{16, 32, 64, 128}	128	2,408,356	0.582
	GaitPTv2-4		{24, 48, 96, 192}	192	5,402,956	1.308
	GaitPTv2-5		{32, 64, 128, 256}	256	9,591,284	2.325
	GaitPTv2-6		{48, 96, 192, 384}	384	21,549,124	5.229
	GaitPTv2-7		{64, 128, 256, 512}	512	38,281,876	9.295
Shallow & Wide	GaitPTv2-1	{2, 2, 4, 1}	{16, 64, 64, 128}	64	1,194,144	0.246
	GaitPTv2-2		{32, 128, 128, 256}	128	4,754,208	0.982
	GaitPTv2-3		{48, 192, 192, 384}	192	10,680,736	2.210
	GaitPTv2-4		{64, 256, 256, 512}	256	18,973,728	3.929
	GaitPTv2-5		{80, 320, 320, 640}	320	29,633,184	6.138
	GaitPTv2-6		{96, 384, 384, 768}	384	42,659,104	8.839
	GaitPTv2-7		{112, 448, 448, 896}	448	58,051,488	12.03
Non-Hierarchical	GaitFormer-1	9	64	64	605,008	0.039
	GaitFormer-2	9	128	128	2,402,688	0.156
	GaitFormer-3	9	192	192	5,393,328	0.351
	GaitFormer-4	9	256	256	9,576,928	0.624
	GaitFormer-5	9	320	320	14,953,488	0.974
	GaitFormer-6	9	384	384	21,523,008	1.403
	GaitFormer-7	9	448	448	29,285,488	1.909

Table 2: Architectural details of the GaitPTv2 and GaitFormer variants. GFLOPs are computed for a forward pass with batch size of 1.

In our experiments, we used slightly modified variants of GaitPT (Catrana, Cosma, and Radoi 2023) and GaitFormer (Cosma and Radoi 2022). GaitPT (Catrana, Cosma, and Radoi 2023) is a hierarchical skeleton transformer with good results for skeleton-based gait recognition on multiple benchmarks. Multiple works in gait recognition (Catrana, Cosma, and Radoi 2023; Fu et al. 2023) have recognized the need of hierarchical processing of skeleton sequences for achieving good performance, by aggregating low-level coordinate information to high level limb movements. We chose GaitPT as a representative architecture for a larger class of models hierarchical pose-based gait models (Fu et al. 2023). In contrast to other architectures (Liao et al. 2020; Fu et al. 2023), GaitPT is a fully attention-based model (Vaswani et al. 2017), which benefits from known scaling properties (Yang et al. 2022; Kaplan et al. 2020; Dehghani et al. 2023) and parallelization of transformer models. In particular, GaitPT is organized similarly to SwinTransformers (Liu et al. 2021), having 4 sequential stages, each stage having spatial transformer layers operating on spatial dimensions of

each skeleton, and a temporal transformer layers aggregating temporal information of the sequence. Readers are referred to the work of Catruna et al. (Catrana, Cosma, and Radoi 2023) for a more detailed description of the model. GaitFormer (Cosma and Radoi 2022) is another full-attention architecture, inspired by vision transformers (Dosovitskiy et al. 2021), in which a simple transformer encoder is used to process sequences of skeletons to output gait representations. In this case, there is no hierarchy of representations, and each skeleton is treated as a single token.

We modify the original GaitPT and GaitFormer implementations to adopt several transformer improvements (Zhang and Sennrich 2019; Wang and Komatsuzaki 2021; Shazeer 2020; Dehghani et al. 2023; Su et al. 2024) for training stabilization and higher throughput without loss of expressive power. In particular, we used "parallel layers" (Wang and Komatsuzaki 2021) by applying the Attention and MLP blocks in parallel, instead of sequentially as in the standard Transformer (Vaswani et al. 2017), we removed bias of QKV projection layers (Dehghani et al. 2023), changed LayerNorm layers to RMS normalization (Zhang and Sennrich 2019), changed the activation function from GeLU (Hendrycks and Gimpel 2016) to SwiGLU (Shazeer 2020) and we used Rotary Positional Embeddings (Su et al. 2024) instead of absolute positional embeddings. For GaitPT, the rotary positional embeddings are instantiated per stage. Furthermore, in the original GaitPT implementation, the final embedding is the direct concatenation of outputs from all 4 stages, resulting in a very large dimensionality – here, we project with a linear layer each stage output to a vector of dimension emb_size . As such, the output concatenation of each stage has dimension $4 \cdot emb_size$, which is projected using another linear layer into dimension emb_size . We name this modified model GaitPTv2. Following the training procedure from SimCLR (Chen et al. 2020), we used a non-linear output projection head during training for both architectures.

Model Configurations Since we are interested in analyzing the impact of different model sizes (measured by number of trainable parameters), we vary only the width, and keep the depth fixed, allowing the use of maximal update parametrization (μP) (Yang et al. 2021) for zero-shot hyperparameter transfer across model widths.

In this work, we analyze two variants of the GaitPTv2 architecture: a deeper but thinner model and a more shallow but wider model (Table 2). Since the original GaitPT is comprised of 4 stages, we vary the GaitPTv2 model size by fixing its depth across each stage and increasing only the transformer model widths and number of heads in terms of a single multiplicative factor. For GaitFormer, we build each model configuration using the corresponding hyperparameters from the third GaitPT stage, as shown in Table 2. For the deeper GaitPTv2, we used 12 layers in Stage 3 and 2 layers everywhere else, inline with SwinTransformers (Liu et al. 2021). For the shallower model, we used only 4 layers in Stage 3, but increased d_{model} for all stages.

Hyperparameters and Pretraining Details

We used μ P parametrization (Yang et al. 2021) to avoid expensive hyperparameter search in larger models. μ P parametrization modifies the learning rate and initializations for feed-forward layers in proportion to the relative width compared to a base model. As such, the optimal learning rate found for a base model size can be directly adapted to larger model scales, as long as the depth of the model remains fixed. In this way, our trained models are not affected by sub-optimal hyperparameter choices and enable us to make a fair representation of model performance across scales. Other works in this area (Xie et al. 2023) fix hyperparameters for all model scales, which is sub-optimal, since it considerably affects results. As far as we know, this work is the first reproduction of μ P in self-supervised gait recognition.

As opposed to scaling analyses in NLP (Kaplan et al. 2020), we do not limit model training on a single epoch, since transformers usually benefit from processing multiple epochs (Muennighoff et al. 2023; Xie et al. 2023). Furthermore, contrastive learning pretraining schemes such as SimCLR (Chen et al. 2020) require multiple training epochs to achieve a reasonable performance due to the increased diversity of data augmentations.

For modeling power-law scaling of model and data size, we train our models for a fixed number of 25 training epochs across data subsets and model scales. We used a fixed batch size of 256 samples across all model scales, which includes the 2 augmented views for each walking sequence. For modeling compute allocation, we increased the batch size to 2048 for all models. All models are trained using AdamW optimizer with mixed-precision, and we used a learning rate of 0.0016 for the smallest model (which is adapted using μ P across model scales) with a cosine learning rate schedule with 1024 iterations for warm-up. We used two NVIDIA H100 / A100 GPUs for training.

Model Evaluation: Controlled and In-the-Wild

For computing model performance across scales, we evaluated the pretrained model in a zero-shot manner (with no fine-tuning) on 4 datasets: CASIA-B (Yu, Tan, and Tan 2006), PsyMo (Cosma and Radoi 2024), GREW (Zhu et al. 2021) and Gait3D (Zheng et al. 2022). For performance evaluation in controlled gait recognition settings we used CASIA-B and PsyMo, and for performance evaluation in realistic scenarios (i.e., "in-the-wild") we used GREW and Gait3D. Both CASIA-B and PsyMo have similar dedicated evaluation procedures that aim to have a fine-grained measure of performance across viewpoints and walking variations. For controlled scenarios, we follow each dataset's evaluation protocol and compute the average performance across viewpoints and scales, excluding identical view cases, and aggregate both metrics in a single value. For in-the-wild performance evaluation we follow each dataset's evaluation procedure and compute rank-5 recognition accuracy and aggregate both metrics in a single value.

Modeling the Power-law for Gait Recognition

We hypothesize that gait recognition follows power-law scaling similar to language modeling (Kaplan et al. 2020; Hoffmann et al. 2022) and other domains (Xie et al. 2023; Cherti et al. 2023). In the case of contrastive learning for gait recognition, we chose to model performance in terms of aggregated downstream accuracy instead of relying on loss value. Consequently, we assume that the aggregated performance $P(N, D)$ in terms of accuracy of a gait recognition model depends on the number of model parameters N and the dataset size D measured in number of skeleton sequences. Following similar works (Kaplan et al. 2020), we assume that $P(N, D) = N^\alpha + D^\beta + E$, where E is a constant irreducible error term. In our experiments, we compute separate scaling laws for model and dataset scale, respectively. If we fix the model size into a fixed set of scales and varied the dataset size, the N^α term becomes constant, and the data scaling law becomes $P_N(D) = D^\beta + E$. Similarly, if we fixed the dataset size, the model scaling law becomes $P_D(N) = N^\alpha + E$. The parameters α and β can be found through least-squares linear regression on a set of model performance values across scales, in a log-log plot.

Results

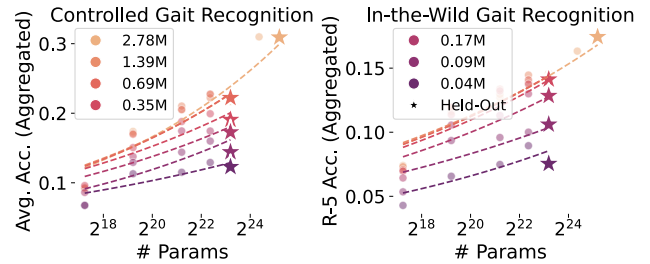


Figure 2: Scaling trends for increasing the model size by parameter count, across multiple dataset sizes. We compute scaling trends only on the data points marked with a "●" symbol, while the "★" data point is used for validation. Increasing the parameter count yields a predictable positive increase in performance.

Power-Law Scaling of GaitPTv2

In this subsection, we present power-law scaling trends and analyze extrapolations across model and data scales. Here, we used the "Deep & Thin" variant of GaitPTv2.

Model scaling for self-supervised gait recognition. In Figure 2 we show model scaling behaviour for several sizes of the pretraining dataset. Larger models have almost always better performance, regardless of the amount of pretraining data. When computing trend lines, we used all but the last data point, and used the final training run (denoted by the ★) for validation. We only train GaitPTv2-7 on the largest data subset, due to the computational constraints of our setup. The largest model's performance (i.e., $\mu = 7$ and $\mu = 6$) closely follows the trend line and their performance can be extrapolated from training smaller scale models.

Model Scaling			Data Scaling		
# Sequences	α (C)	α (W)	# Params	β (C)	β (W)
43.4K (1.56%)	0.100	0.118	0.15M ($\mu = 1$)	0.070	0.094
86.9K (3.12%)	0.139	0.099	0.61M ($\mu = 2$)	0.078	0.100
173.7K (6.25%)	0.122	0.111	2.41M ($\mu = 3$)	<u>0.109</u>	<u>0.102</u>
347.5K (12.5%)	0.124	<u>0.113</u>	5.4M ($\mu = 4$)	0.102	0.091
694.9K (25.0%)	<u>0.148</u>	0.113	9.59M ($\mu = 5$)	0.114	0.108
1389.9K (50.0%)	0.145	0.109			
2779.8K (100.0%)	0.164	0.112			

Table 3: Scaling parameters α and β for both model size and dataset size in zero-shot controlled and in-the-wild gait recognition scenarios.

In Table 3 on the left-hand side, we show model scaling parameters α for multiple subsets of our dataset. Even though training on the largest subset has the steepest scaling parameter for controlled scenarios, for in-the-wild settings parameters are fairly close to one another. This is due to the fact that the models are likely undertrained and could significantly benefit from more training. Training on smaller amounts of data usually results in overfitting, and by observing good results for smaller models in this scenario indicates that the models are not trained to saturation.

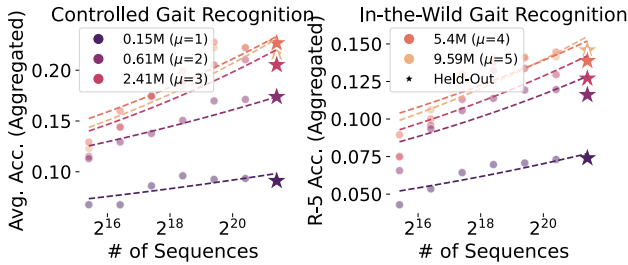


Figure 3: Scaling trends for increasing the dataset size in terms of the number of skeleton sequences, across multiple model sizes. We compute scaling curves from points marked "•". The point marked with "★" is used for validation.

Data scaling for self-supervised gait recognition. In Figure 3 we show data scaling behaviour for several model sizes when trained with progressively larger dataset sizes. All model scales benefit from increasing the size of the pre-training dataset in both controlled and in the wild scenarios. When computing the trend line, we used all but the largest data scale, and the final training run (denoted by the ★) for validation. While the final point is close to the trend line, there seems to be a saturation point at larger amounts of data, likely due to the added noise in the dataset. We explore the effect of data quality on scaling below. We could argue that gait has comparatively less entropy than natural language, where scaling has been more extensively studied (Kaplan et al. 2020), which might lead to faster data saturation (Caballero et al. 2022). As opposed to collecting text data (Villalobos et al. 2022), diverse gait data is easier to gather from video streams, as different environments may lead to radically different ways of walking. In Table 3, on the right-hand side, we show the data scaling parameters β for multiple model sizes. From our experiments, the larger the

model size, the more it is able to consume data for this task. The model is likely under-trained and could significantly improve its performance with training to saturation.

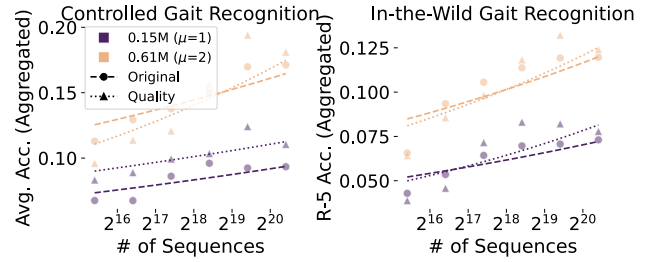


Figure 4: Comparison of data scaling behaviour between models trained on high quality samples versus models trained on samples from the original set.

Effect of skeleton data quality To study the effect of modifying the data quality on scaling trends, we propose a simple heuristic to order skeleton sequences in terms of the quality of extracted human poses. Considering that a 2D skeleton sequence S is comprised of a set of $J = 18$ joints having 2 coordinate values (x_j, y_j) and a confidence score c_j , we compute, for each joint, the average confidence \bar{c}_j and the variance of the confidence σ_{c_j} across sequence length.

The quality score is defined as $Q_S = \sum_{j=0}^J (\bar{c}_j - \log \sigma_{c_j})$. Assuming high quality sequences should have high average confidence and low confidence variance across time, ordering the dataset by Q_S gives a monotonically increasing set of skeleton sequences by quality of extraction. As such, in each subset we sample the top quality sequences. Figure 4 shows the data scaling properties of training with high quality samples compared to the original, randomly sampled subset. Training with high quality samples yields consistent better performance and a slightly steeper scaling curve. For example, for $\mu = 2$, the scaling parameter β for the high quality samples is 0.1319 versus 0.078 on controlled scenarios, and 0.125 versus 0.099 in in-the-wild scenarios.

Scaling the compute budget across architectures

In this subsection, we present our analyses in terms of amount of compute (FLOPs), and provide a comparison between GaitPTv2 and GaitFormer. Here, we used the "Shallow & Wide" variant of GaitPTv2.

In Figure 5 we show training IsoFLOPs curves for GaitPTv2 and GaitFormer. We fixed compute budgets and selected the closest model that reached that budget during training. Each curve was obtained by fitting a linear model to model accuracy as function of number of parameters: $P(\alpha) = \alpha \log_2(N) + E$, where P is downstream controlled gait recognition accuracy, N is the number of parameters and E is the intercept. In our scenario, *smaller models have a more efficient use of compute given enough data*. Similarly, in Figure 6, we plot training IsoFLOPs curves obtained by fitting a linear model in the form $P(\beta) = \beta \log_2(D) + E$, where D is the number of training gait sequences. In this case, *increasing the number of gait sequences seen during*

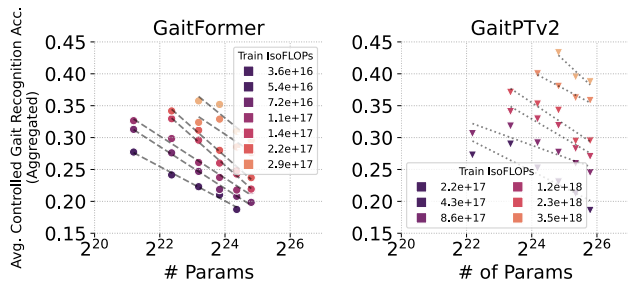


Figure 5: Training IsoFLOPs curves for GaitPTv2 and GaitFormer, comparing number of parameters and controlled gait recognition accuracy. The points on each curve utilize approximately the same amount of training compute.

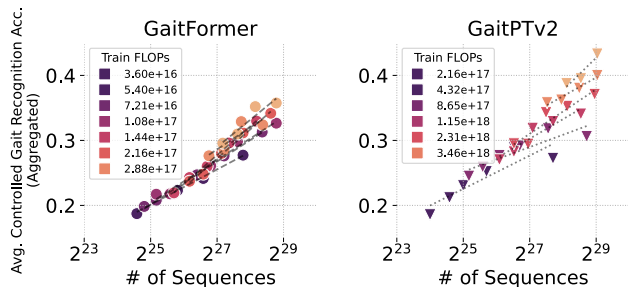


Figure 6: Training IsoFLOPs curves for GaitPTv2 and GaitFormer, comparing number of training sequences and controlled gait recognition accuracy. The points on each curve utilize approximately the same amount of training compute.

training improves performance across compute budgets.

In Figure 7 we show Iso-Accuracy curves for GaitPTv2 and Gaitformer as well as Iso-FLOP contours. We approximate the amount of FLOPs consumed by each model as $6.5 \cdot ND$ for GaitPTv2 and $\sim 2 \cdot ND$ for GaitFormer, by fitting a linear model of the form $C(\gamma) = \gamma ND$. Based on these results, GaitFormer, a simple transformer encoder model, is more efficient in terms of compute for the same number of parameters compared to GaitPTv2. However, *GaitPTv2 obtains better accuracy because it uses more compute*. The main difference stems from the way skeleton sequences are processed between the two models. In the case of GaitFormer, each "token" is considered a flattened skeleton, which discards explicit spatial relationship between joints. GaitPTv2, however, processes sequences hierarchically (Catruna, Cosma, and Radoi 2023; Liu et al. 2021), from single joints to body parts, having a larger effective context length, and, in turn, more compute expenditure per sequence. As a consequence, *GaitPTv2 scales better with amount of data compared to GaitFormer*. We show this result in Figure 8: we select the most efficient models in terms of compute (from Figures 5 and 6) and plot trends across parameter counts and number of training sequences. Scaling parameter count does not show a substantial difference between models, but the gap is evident when scaling number of training gait sequences: GaitPTv2 is using more compute

per gait sequence, resulting in better downstream gait recognition performance.

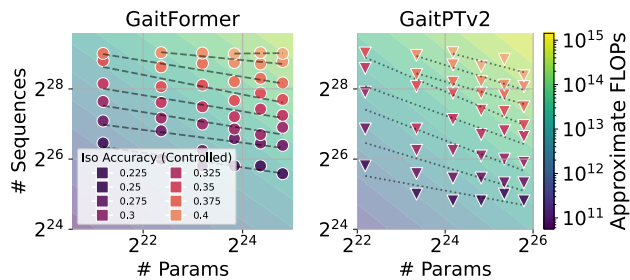


Figure 7: Iso-Accuracy curves. The background is colored using IsoFLOPs contours for both models.

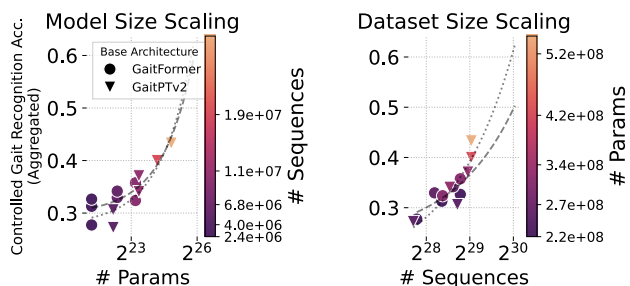


Figure 8: Trends for using the most efficient models across FLOP budgets.

Conclusions

Our study represents the first attempt at constructing scaling laws for self-supervised gait recognition, exploring the dynamics of model performance, data quantity, and computational resources to downstream zero-shot performance on both controlled gait recognition scenarios and in-the-wild scenarios. We gathered a dataset of 2.7M skeleton sequences in-the-wild, the largest reported in literature, and used an improved version of GaitPT for pretraining. We showed that, power-laws performance trends do apply to gait recognition, enabling practitioners to predict the performance of models by training only smaller versions of models, and on smaller amounts of data. We also presented the first reproduction of μP in gait recognition, which allowed us to directly transfer hyperparameters from small models to larger ones without additional search.

We further isolate architectural contributions by comparing GaitPTv2 with GaitFormer under controlled compute budgets and showed that GaitPTv2 scales better with data because, by its hierarchical design, it is able to use more compute per gait sequence. Our work represents a promising avenue for further research into scaling self-supervised gait recognition.

Acknowledgments

This research was supported by the project “Romanian Hub for Artificial Intelligence - HRIA”, Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MyS-MIS no. 351416.

References

- Alabdulmohsin, I. M.; Zhai, X.; Kolesnikov, A.; and Beyer, L. 2024. Getting vit in shape: Scaling laws for compute-optimal model design. *Advances in Neural Information Processing Systems*, 36.
- Bansal, Y.; Ghorbani, B.; Garg, A.; Zhang, B.; Cherry, C.; Neyshabur, B.; and Firat, O. 2022. Data scaling laws in NMT: The effect of noise and architecture. In *International Conference on Machine Learning*, 1466–1482. PMLR.
- Caballero, E.; Gupta, K.; Rish, I.; and Krueger, D. 2022. Broken neural scaling laws. *arXiv preprint arXiv:2210.14891*.
- Caballero, E.; Gupta, K.; Rish, I.; and Krueger, D. 2023. Broken Neural Scaling Laws. In *The Eleventh International Conference on Learning Representations*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33: 9912–9924.
- Catrana, A.; Cosma, A.; and Radoi, E. 2023. GaitPT: Skeletons Are All You Need For Gait Recognition. *arXiv preprint arXiv:2308.10623*.
- Chao, H.; Wang, K.; He, Y.; Zhang, J.; and Feng, J. 2021. GaitSet: Cross-view Gait Recognition through Utilizing Gait as a Deep Set. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Cherti, M.; Beaumont, R.; Wightman, R.; Wortsman, M.; Ilharco, G.; Gordon, C.; Schuhmann, C.; Schmidt, L.; and Jitsev, J. 2023. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2818–2829.
- Cortes, C.; Jackel, L. D.; Solla, S.; Vapnik, V.; and Denker, J. 1993. Learning curves: Asymptotic values and rate of convergence. *Advances in neural information processing systems*, 6.
- Cosma, A.; and Radoi, E. 2022. Learning Gait Representations with Noisy Multi-Task Learning. *Sensors*, 22(18).
- Cosma, A.; and Radoi, E. 2023. GaitMorph: Transforming Gait by Optimally Transporting Discrete Codes. *arXiv:2307.14713*.
- Cosma, A.; and Radoi, E. 2024. PsyMo: A Dataset for Estimating Self-Reported Psychological Traits from Gait. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4603–4613.
- Cosma, A.; and Radoi, I. E. 2021. WildGait: Learning Gait Representations from Raw Surveillance Streams. *Sensors*, 21(24): 8387.
- Cătrună, A.; Cosma, A.; and Rădoi, E. 2024. The Paradox of Motion: Evidence for Spurious Correlations in Skeleton-based Gait Recognition Models. *arXiv:2402.08320*.
- Dehghani, M.; Djolonga, J.; Mustafa, B.; Padlewski, P.; Heek, J.; Gilmer, J.; Steiner, A. P.; Caron, M.; Geirhos, R.; Alabdulmohsin, I.; et al. 2023. Scaling vision transformers to 22 billion parameters. In *International Conference on Machine Learning*, 7480–7512. PMLR.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Fan, C.; Hou, S.; Wang, J.; Huang, Y.; and Yu, S. 2023. Learning Gait Representation From Massive Unlabelled Walking Videos: A Benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12): 14920–14937.
- Fan, C.; Ma, J.; Jin, D.; Shen, C.; and Yu, S. 2024. Skeleton-gait: Gait recognition using skeleton maps. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 1662–1669.
- Fang, H.-S.; Li, J.; Tang, H.; Xu, C.; Zhu, H.; Xiu, Y.; Li, Y.-L.; and Lu, C. 2022. AlphaPose: Whole-Body Regional Multi-Person Pose Estimation and Tracking in Real-Time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Fu, Y.; Meng, S.; Hou, S.; Hu, X.; and Huang, Y. 2023. Gp-gait: Generalized pose-based gait recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19595–19604.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hernandez, D.; Kaplan, J.; Henighan, T.; and McCandlish, S. 2021. Scaling laws for transfer. *arXiv preprint arXiv:2102.01293*.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L. A.; Welbl, J.; Clark, A.; Hennigan, T.; Noland, E.; Millican, K.; van den Driessche, G.; Damoc, B.; Guy, A.; Osindero, S.; Simonyan, K.; Elsen, E.; Vinyals, O.; Rae, J.; and Sifre, L. 2022. An empirical analysis of compute-optimal large language model training. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 30016–30030. Curran Associates, Inc.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Li, X.; Makihara, Y.; Xu, C.; Yagi, Y.; Yu, S.; and Ren, M. 2020. End-to-end model-based gait recognition. In *Proceedings of the Asian conference on computer vision*.

- Liao, R.; Yu, S.; An, W.; and Huang, Y. 2020. A model-based gait recognition method with body pose and human prior knowledge. *Pattern Recognition*, 98: 107069.
- Lin, B.; Zhang, S.; Wang, M.; Li, L.; and Yu, X. 2022. Gaitgl: Learning discriminative global-local feature representations for gait recognition. *arXiv preprint arXiv:2208.01380*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *arXiv:2103.14030*.
- Makihara, Y.; Mannami, H.; Tsuji, A.; Hossain, M.; Sugiura, K.; Mori, A.; and Yagi, Y. 2012. The OU-ISIR Gait Database Comprising the Treadmill Dataset. *IPSJ Trans. on Computer Vision and Applications*, 4: 53–62.
- Muennighoff, N.; Rush, A. M.; Barak, B.; Scao, T. L.; Tazi, N.; Piktus, A.; Pyysalo, S.; Wolf, T.; and Raffel, C. 2023. Scaling Data-Constrained Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Nasseri, M. H.; Moradi, H.; Hosseini, R.; and Babae, M. 2021. Simple online and real-time tracking with occlusion handling. *arXiv preprint arXiv:2103.04147*.
- Nixon, M. S.; Tan, T. N.; and Chellappa, R. 2005. *Human Identification Based on Gait (The Kluwer International Series on Biometrics)*. Berlin, Heidelberg: Springer-Verlag. ISBN 0387244247.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Rosenfeld, J. S.; Rosenfeld, A.; Belinkov, Y.; and Shavit, N. 2019. A constructive prediction of the generalization error across scales. *arXiv preprint arXiv:1909.12673*.
- Sablayrolles, A.; Douze, M.; Schmid, C.; and Jégou, H. 2019. Spreading vectors for similarity search. In *International Conference on Learning Representations*.
- Shazeer, N. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.
- Shin, K.; Kwak, H.; Kim, S. Y.; Ramström, M. N.; Jeong, J.; Ha, J.-W.; and Kim, K.-M. 2023. Scaling law for recommendation models: Towards general-purpose user representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4596–4604.
- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.
- Sutton, R. 2019. The bitter lesson. *Incomplete Ideas (blog)*, 13(1).
- Teepe, T.; Khan, A.; Gilg, J.; Herzog, F.; Hörmann, S.; and Rigoll, G. 2021. GaitGraph: Graph Convolutional Network for Skeleton-Based Gait Recognition. In *2021 IEEE International Conference on Image Processing (ICIP)*, 2314–2318.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Villalobos, P.; Sevilla, J.; Heim, L.; Besiroglu, T.; Hobbhahn, M.; and Ho, A. 2022. Will we run out of data? An analysis of the limits of scaling datasets in Machine Learning. *arXiv preprint arXiv:2211.04325*.
- Wang, B.; and Komatsuzaki, A. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Wei, Y.; Dai, Q.; and Hu, H. 2023. On data scaling in masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10365–10374.
- Xu, Y.; Zhang, J.; Zhang, Q.; and Tao, D. 2022. ViTPose: Simple Vision Transformer Baselines for Human Pose Estimation. In *Advances in Neural Information Processing Systems*.
- Yang, G.; Hu, E.; Babuschkin, I.; Sidor, S.; Liu, X.; Farhi, D.; Ryder, N.; Pachocki, J.; Chen, W.; and Gao, J. 2021. Tuning large neural networks via zero-shot hyperparameter transfer. *Advances in Neural Information Processing Systems*, 34: 17084–17097.
- Yang, G.; Hu, E. J.; Babuschkin, I.; Sidor, S.; Liu, X.; Farhi, D.; Ryder, N.; Pachocki, J.; Chen, W.; and Gao, J. 2022. Tensor Programs V: Tuning Large Neural Networks via Zero-Shot Hyperparameter Transfer. *arXiv:2203.03466*.
- Yao, Y.; and Wang, Y. 2023. Research without Re-search: Maximal Update Parametrization Yields Accurate Loss Prediction across Scales. *CoRR*, abs/2304.06875.
- Yu, S.; Tan, D.; and Tan, T. 2006. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 4, 441–444. IEEE.
- Zhang, B.; and Sennrich, R. 2019. Root Mean Square Layer Normalization. In *Advances in Neural Information Processing Systems 32*. Vancouver, Canada.
- Zhang, Z.; Tran, L.; Liu, F.; and Liu, X. 2019. On Learning Disentangled Representations for Gait Recognition. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Sep. 2019.
- Zheng, J.; Liu, X.; Liu, W.; He, L.; Yan, C.; and Mei, T. 2022. Gait Recognition in the Wild with Dense 3D Representations and A Benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20228–20237.
- Zhu, Z.; Guo, X.; Yang, T.; Huang, J.; Deng, J.; Huang, G.; Du, D.; Lu, J.; and Zhou, J. 2021. Gait Recognition in the Wild: A Benchmark. In *IEEE International Conference on Computer Vision (ICCV)*.
- Zou, S.; Fan, C.; Xiong, J.; Shen, C.; Yu, S.; and Tang, J. 2024. Cross-Covariate Gait Recognition: A Benchmark. In *Proceedings of the AAAI Conference on Artificial Intelligence*.