

Skeleton-in-Context: Unified Skeleton Sequence Modeling with In-Context Learning

Xinshun Wang^{1,2,*}, Zhongbin Fang^{1,*}, Xia Li³, Xiangtai Li⁴, Chen Chen⁵, Mengyuan Liu², ✉

¹Sun Yat-sen University

²National Key Laboratory of General Artificial Intelligence, Peking University, Shenzhen Graduate School

³Department of Computer Science, ETH Zurich ⁴S-Lab, Nanyang Technological University

⁵Center for Research in Computer Vision, University of Central Florida

{wangxsh36, fangzhh5}@mail2.sysu.edu.cn, xia.li@inf.ethz.ch, xiangtai94@gmail.com

chen.chen@crcv.ucf.edu, liumengyuan@pku.edu.cn

<https://github.com/fanglaosi/Skeleton-in-Context>

Abstract

In-context learning provides a new perspective for multi-task modeling for vision and NLP. Under this setting, the model can perceive tasks from prompts and accomplish them without any extra task-specific head predictions or model fine-tuning. However, skeleton sequence modeling via in-context learning remains unexplored. Directly applying existing in-context models from other areas onto skeleton sequences fails due to the similarity between inter-frame and cross-task poses, which makes it exceptionally hard to perceive the task correctly from a subtle context. To address this challenge, we propose Skeleton-in-Context (SiC), an effective framework for in-context skeleton sequence modeling. Our SiC is able to handle multiple skeleton-based tasks simultaneously after a single training process and accomplish each task from context according to the given prompt. It can further generalize to new, unseen tasks according to customized prompts. To facilitate context perception, we additionally propose a task-unified prompt, which adaptively learns tasks of different natures, such as partial joint-level generation, sequence-level prediction, or 2D-to-3D motion prediction. We conduct extensive experiments to evaluate the effectiveness of our SiC on multiple tasks, including motion prediction, pose estimation, joint completion, and future pose estimation. We also evaluate its generalization capability on unseen tasks such as motion-in-between. These experiments show that our model achieves state-of-the-art multi-task performance and even outperforms single-task methods on certain tasks.

1. Introduction

The idea of multi-task models has attracted much research interest over the years in a wide range of areas [6, 16, 17, 58], such as computer vision and natural language processing

* Equal contribution. ✉ Corresponding author is Mengyuan Liu.

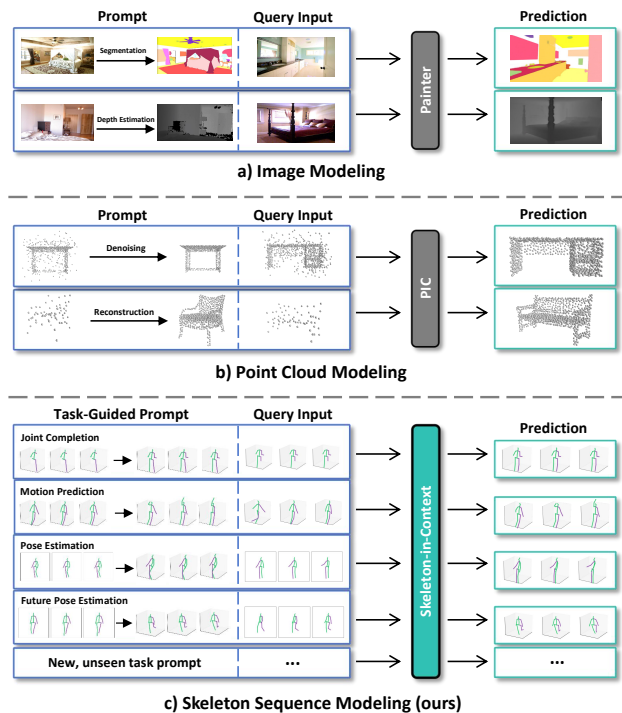


Figure 1. In-context learning in a) image modeling [58], b) point cloud modeling [16], and c) skeleton sequence modeling (ours).

(NLP). To build such models, most previous works focus on deploying task-specific heads on a shared backbone and learning them through pretraining and fine-tuning. Models with such task-specific architectures are not learned in an end-to-end manner and are hard to generalize to new, unseen tasks and data. Most existing methods for skeleton-based human-centric tasks rely on task-specific designs and/or non-end-to-end frameworks to accomplish multiple tasks. Several works [9, 36] learn unified human motion representations. In particular, MotionBERT [77] adopts two-stage training, including pre-training on the 2D-3D lifting task

and fine-tuning on each task. UniHCP [12] connects a well-designed shared backbone network to specific task heads according to different tasks. Meanwhile, UPS [17] unifies action labels and motion sequences into text-based language sequences and fine-tunes them on each task for better performance. These models can have a limited scope of tasks they are capable of handling and lack flexibility. This limits their capability to generalize to unseen tasks and data. To overcome these challenges, in-context learning represents a new trend to build generalist models, which was originally proposed in NLP [6, 45, 46] and then introduced into other areas such as images and point clouds. Some generalist models [5, 58, 59, 72] utilize the Masked Image Modeling (MIM) framework to exhibit multi-task capability via in-context learning in image processing, as shown in Fig. 1(a). Similarly, in Fig. 1(b), Point-In-Context [16] proposes a key joint sampling module and introduces in-context learning into point cloud processing. Specifically, in-context models learn patterns of provided input-output pairs, called in-context examples or prompts, and then perform the desired task according to the prompt. An in-context model can accomplish diverse tasks end-to-end without any task-specific heads or fine-tuning. It also has excellent generalization capability. Directly adapting existing in-context models from other areas onto skeleton sequences fails because they largely focus on static objects (images or point clouds) and thus cannot capture spatial-temporal dependencies in skeleton sequences. To the best of our knowledge, in-context learning remains unexplored in 3D skeleton sequence modeling.

In this paper, we propose Skeleton-in-Context (SiC), a novel framework for skeleton sequence modeling, which is able to process multiple tasks simultaneously after a single training process without any task-specific designs, as Fig. 1(c) shows. Specifically, SiC perceives tasks from context, i.e., the prompt from a desired task provided as additional input, and then accomplishes the task for the query. SiC employs two types of prompts to facilitate context perception: the task-guided prompt (TGP) and the task-unified prompt (TUP). TGP provides a task template for the model, giving the model the ability to make analogies. This is the key to the model’s capability to generalize to new tasks, as we can guide the model to do unseen tasks by customizing specific prompts. TUP can be implemented in two ways: static and dynamic. The static TUP leverages prior knowledge shared among tasks, while the dynamic TUP adaptively learns to perceive task context. Skeleton-in-Context provides a new perspective for training a skeleton-based multi-task in-context model. Since there is no existing benchmark for measuring such a model, we establish a new in-context benchmark based on four tasks, including three skeleton-based tasks: motion prediction, pose estimation, and joint completion, and a new task: future pose estimation.

Our main contributions are as follows: 1) We introduce

a novel framework, Skeleton-in-Context (SiC), designed to handle multiple skeleton-based tasks simultaneously. Our approach learns the potential pattern of the given prompt and performs it on the query sample. 2) We establish a new 3D human in-context learning benchmark comprised of multiple skeleton-based tasks, such as motion prediction, pose estimation, joint completion, and future motion estimation. We further benchmark several representative in-context learning baselines for reference. 3) We conduct extensive experiments to evaluate our SiC. SiC achieves state-of-the-art results in multi-tasking and even outperforms task-specific models on certain tasks. Moreover, SiC also generalizes well to unseen tasks such as motion in-between.

2. Related Work

Skeleton-Based Motion Analysis. Many of the skeleton-based human-centric tasks have shared similar development patterns in terms of the methods used to address them. On human motion prediction, early successes are achieved with RNNs and CNNs [28, 40]. Nowadays, the task is dominated by Graph Convolutional Networks (GCNs). LTD [38] considers each joint coordinate as a node in a graph on which to employ learnable spatial graph convolutions. Follow-up works [13, 39, 60, 61] focus on designing better graph convolution networks. Different from motion prediction, the goal of motion synthesis is to complete the human motion sequence. To complete the missing parts of human motion, many works adopt convolutional models [22, 26, 68] or generative adversarial networks [7, 20, 23, 76]. One category of 3D pose estimation [18, 32, 34] is to lift the estimated 2D pose to 3D, which often employs spatial-temporal convolution [8, 10, 11, 41, 44, 57, 67] or transformer-based [14, 49, 71, 73–75] methods. However, these approaches focus on designing customized architectures specific to a task, which are limited in range of applications. To this end, we propose to unify the input and output formats of multiple tasks and utilize a single model to process multiple tasks simultaneously via in-context learning.

In-Context Learning. The emergence of in-context learning provides a new paradigm for achieving multi-tasking, which originated from Natural Language Processing (NLP) [1, 15, 45–47]. It endows models with the capability to handle multiple tasks and even unseen tasks by incorporating domain-specific input-target pairs, known as prompts, into the query sample [6]. Following NLP, in-context learning is extended to computer vision in some work, such as 2D images [51, 56, 58, 59, 62, 72] and 3D point clouds [3, 16, 64]. Specifically, Visual Prompt [5], Painter [58], and Point-In-Context [16] are pioneers of in-context methods in computer vision. They verify the effectiveness and flexibility of in-context learning in vision-based tasks, and most of them adopt pure-transformer-based approaches [4, 21, 29, 63]. Meanwhile, several approaches

deeply explore Prompt Engineering [3, 51, 72] and awe-inspiring generalization [16, 59, 64] in in-context learning. Recently, LVM [2] unifies most vision tasks and data via joint co-training on one visual in-context model. Compared to the works focused on *static* objects (images and point clouds), our work mainly delves into modeling in-context *dynamic* skeleton sequences.

Unified Skeleton-Centric Models. A generalist model that can excel in handling multiple tasks is a profound step toward general artificial intelligence [45, 58]. Previous works [25, 27, 31, 35, 50, 53, 66, 69] give the effort to explore co-training several human-centric tasks. Meanwhile, some works [30, 42] show great progress in simultaneously predicting parsing and pose. Recent research [9, 36] demonstrates the effectiveness of modeling unified human motion sequences. In particular, UniHCP [12] designs specific task heads according to different 2D human-centric tasks. Motion-BERT [77] is pre-trained in the 2D-3D lifting task and fine-tuned on each task. UPS [17] unifies heterogeneous human behavior understanding tasks via language sequences and cycle training for each task in order. Nevertheless, these works are limited by the dilemma of complex training stages or task heads, which introduce extra costs. We integrate the output format of different tasks in skeleton sequence format and train the model just once.

3. In-Context Skeleton Sequence Modeling

3.1. Formulation

Since the skeleton sequence data in skeleton-based tasks are sequential by nature, we treat the skeleton at each frame as a token, which can be processed by transformer [54]. Following previous works [5, 16, 58], we curate an in-context skeleton sequence paradigm. Given an input-target skeleton sequence pair as a task-guided prompt $G_i^k = \{I_i^k, T_i^k\}$ and a query input Q_j^k , where $I_i^k, T_i^k, Q_j^k \in \mathbb{R}^{F \times J \times 3}$ each have F frames and J joints, the model is expected to perceive the context from the prompt, i.e., which task it represents, and perform the corresponding operations on the query in analogy to the prompt. This process can be formulated as:

$$y = f(G_i^k, Q_j^k), \quad (1)$$

where k represents the task index, while i and j indicate different example indexes. Thus, we can determine what task will be performed on the query sample by selecting the task-specific prompt from the corresponding training set.

3.2. Task Definitions

To measure multitasking capability, we establish a new in-context benchmark from common skeleton-based datasets, including H3.6M [24], AMASS [37], and 3DPW [55]. In this benchmark, we select three common tasks: motion prediction, pose estimation, and joint completion, and a new

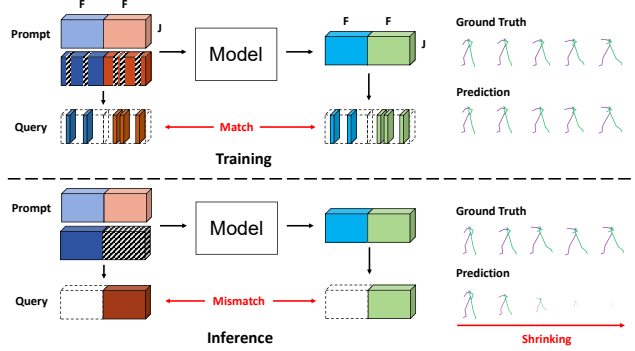


Figure 2. *Top*: Training with MIM-style framework [21, 43] used in previous works [5, 16, 58]. The model is able to reconstruct the masked frames well. *Bottom*: During inference, the reconstructed sequence gradually shrinks the skeleton as time goes by because the model only learns frame interpolation during training. Once the model loses subsequent reference frames when interpolating frames, the generated sequence will tend to shrink.

task: future pose estimation. Consequently, we conduct extensive experiments on a large and diverse dataset collection, which consists of 176K skeleton sequence samples.

Motion Prediction. Motion prediction [38] is to predict future motion based on historical motion. There is continuity between the input and output sequences, which constitute a complete motion. Thus, the input I and target T are complete skeleton sequences with shapes of $(F, J, 3)$.

Pose Estimation. Pose estimation [18, 32, 34] is to lift 2D poses to 3D. The model takes the detected 2D skeleton sequence $I \in \mathbb{R}^{F \times J \times 2}$, which contain noise, as input, and estimates the accurate 3D skeleton sequence $T \in \mathbb{R}^{F \times J \times 3}$. Note that we expand the 2D input to 3D by filling the 3rd dimension with 0, in order to unify the data format with other tasks. Therefore, the input skeleton sequence becomes the same shape $(F, J, 3)$ as the target sequence.

Future Pose Estimation. Based on prediction and estimation, we propose a new task: future pose estimation, which is to predict future 3D motion based on history 2D poses. As 2D images are easier to obtain than 3D data, this task unlocks a host of new real-world applications. Given a historical 2D skeleton sequence $I \in \mathbb{R}^{F \times J \times 3}$ (one dimension with 0), future pose estimation expects the model to output its future 3D skeleton sequence $T \in \mathbb{R}^{F \times J \times 3}$. To our knowledge, no work has ever proposed future pose estimation.

Joint Completion. Joint completion is to recover skeleton sequences that have missing joints. This task can reflect cases where joints may be missing or displaced when collected with Lidar cameras [33, 48]. We set two levels to verify the ability to complete joints, which are 40% and 60% of the joints missing in each frame. Also, the coordinates of missing joints are replaced with 0, with the goal of unifying the input skeleton sequence format with $(F, J, 3)$.




Figure 3. Overall framework of our Skeleton-in-Context. Specifically, we establish a skeleton bank by integrating training sets under different tasks, which contain a large amount of input-target pairs performing different tasks. Next, we randomly select a sample pair as the task-guided prompt (TGP) and a query input from the skeleton bank, undergo encoding and concatenating, respectively, and then input them into the transformer in parallel. In particular, during this process, the query input and task-unified prompt (TUP) are combined to form a new query. After iterating n_1 times, the TGP and query are aggregated through $\mathcal{A}(\cdot)$ and then input into the transformer for n_2 iterations. Lastly, the second half of the model output is used as our prediction.

4. Skeleton-in-Context

In-context learning on other data modalities such as images and point clouds [5, 16, 58] largely employ Mask Image Modeling (MIM) [21] to reconstruct the randomly masked elements. During inference, only the target elements are masked. On skeleton sequences, MIM fails due to spatial-temporal complexity and inter-pose similarity, which makes it outstandingly hard to perceive the task correctly from a subtle context. In our Skeleton-in Context, we employ a new framework to address this challenge.

4.1. Skeleton Sequence Prompt

Task-Guided Prompt. In in-context learning, a prompt is needed for the model to learn the underlying task pattern from context. As mentioned in Sec. 3.1, we propose the Task-Guided Prompt (TGP) $G_i^k = \{I_i^k, T_i^k\}$, which comprises of an example skeleton pair performing a specific task k . In this way, our proposed model is able to perform a variety of tasks according to the provided prompts that could be seen or even unseen in the training set. It is worth noting that our model can perform unseen tasks and transfer to datasets that are different styles from the training dataset. It will be elaborated in Sec. 5.2.

Avoiding Over-Fitting. Previous works have made great progress in unifying different tasks via the MIM-style framework [5, 16, 58]. Compared with static objects in 2D and 3D, such as images and point clouds, the obstacle to in-context learning of skeleton sequences lies in the wondrous similarity of each frame, so directly using the random masking strategy [5, 16, 58] will cause the model to fall into over-fitting, leading it to only learn simple interpolation frames rather than capturing meaningful motion sequence representations during the training stage. As shown in Fig. 2, when

we directly utilize the MPM-style [21, 43] framework to train our model, although it can perfectly reconstruct the masked frame during training, the experimental results reveal that it outputs a devastating result during inference. It is obvious that the output skeleton sequence gradually shrinks as the length increases because the next frame of the current frame is 0, and interpolating with 0 will lead to shrinking. Such a phenomenon underscores the challenges of learning contextual information about dynamic skeleton sequences. **Task-Unified Prompt.** With the Task-Guided Prompt, which exhibits unique capabilities specific to in-domain or out-of-domain tasks, mitigating the gap between different tasks presents a new challenge. Meanwhile, as discussed above, a straight extension of previous works will lead to a simple model that can only perform the interpolation task based on the given skeleton sequence. To tackle these challenges, we introduce the Task-Unified Prompt (TUP), a well-designed and task-agnostic module, which can adaptively learn to incorporate multiple tasks into the framework of in-context learning, i.e., a one-off end-to-end framework for all tasks without having to deploy task-specific heads or fine-tuning. We introduce two versions of TUP:

Version 1 : Static-TUP with static pseudo pose:

$$U^S = \mathcal{E} \left(\frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K (T_m^k) \in \mathbb{R}^{F \times J \times 1} \right), \quad (2)$$

where $\mathcal{E}(\cdot)$ is an embedding layer. M and K indicate the total number of training samples and tasks, respectively.

Version 2 : Dynamic-TUP with dynamic pseudo pose:

$$U^D = W_1 W_2, \quad (3)$$

where $W_1 \in \mathbb{R}^{F \times J \times 1}$ and $W_2 \in \mathbb{R}^{1 \times 1 \times C}$.

Table 1. Comparison of task-specific models, multi-task models, and in-context models on four human skeleton tasks. For Motion Prediction (MP.), Joint Completion (JC.), and Future Pose Estimation (FPE.), we report Mean Per Joint Position Error (mm) [24]. For Pose Estimation (PE.), we additionally report another indicator, Normalized Mean Per Joint Position Error (N-MPJPE) [77]. For joint completion, we report two settings: 40% and 60% of the joints missing in each frame. † represents the multi-task-capable models re-structured from task-specific models. ‡ represents the multi-task-capable models re-structured from multi-stage models. Static-SiC and Dynamic-SiC represent the static and dynamic versions of TUP, respectively.

Methods	Venues	MP. (AMASS)						PE. (H3.6M)		JC. (3DPW)			FPE. (H3.6M)		
		MPJPE ↓						MPJPE ↓	N-MPJPE ↓	MPJPE ↓			MPJPE ↓		
		80ms	160ms	200ms	320ms	400ms	Avg.	Avg.	Avg.	40%	60%	Avg.	200ms	300ms	Avg.
Task-Specific Model: one architecture for one task															
siMLPe [19]	WACV'23	12.1	21.6	25.4	34.4	39.9	26.7	58.9	49.7	48.4	58.5	53.5	75.4	77.6	76.5
EqMotion [65]	CVPR'23	14.4	26.6	31.7	44.0	50.0	33.3	63.2	59.3	40.4	49.7	45.1	69.5	70.8	70.1
STCFormer [52]	CVPR'23	19.5	25.1	28.2	35.8	41.5	28.7	54.6	45.2	31.2	46.7	38.9	60.9	81.8	71.3
GLA-GCN [70]	ICCV'23	19.5	25.1	28.2	35.3	42.0	30.0	49.6	40.4	62.2	72.5	67.4	49.5	51.4	50.5
MotionBERT [77]	ICCV'23	12.9	21.4	25.5	36.3	42.3	26.7	49.4	41.1	29.3	44.3	36.8	75.5	100.3	87.9
Multi-Task Model: one architecture (w. multiple task-specific heads) for multiple tasks															
siMLPe† [41]	WACV'23	19.0	30.0	35.3	46.8	52.4	36.7	74.6	60.0	59.0	67.4	63.2	64.4	68.4	66.4
EqMotion† [65]	CVPR'23	14.0	25.3	30.2	40.5	46.6	31.3	163.3	110.1	38.9	56.0	47.5	88.1	89.1	88.6
STCFormer† [52]	CVPR'23	17.4	28.0	31.5	43.0	40.4	33.9	55.6	45.6	33.6	50.6	42.1	62.2	71.7	67.0
GLA-GCN† [70]	ICCV'23	35.3	39.8	42.3	49.2	56.1	44.5	78.6	62.9	74.1	81.9	78.0	60.6	68.1	64.4
MotionBERT‡ [77]	ICCV'23	14.3	22.8	26.7	36.4	41.7	28.4	55.8	45.7	31.3	46.7	39.0	75.5	96.8	86.2
In-Context Model: one task-agnostic architecture for multiple tasks															
Copy [5, 16]	-	134.7	133.7	133.3	132.1	131.4	133.0	351.0	171.6	196.5	195.5	196.0	316.8	316.6	316.7
PointMAE [43]	ECCV'22	19.1	31.2	36.6	54.9	67.7	41.9	144.3	96.41	67.6	79.0	73.3	139.9	142.9	141.4
PIC [16]	NeurIPS'23	25.0	40.2	46.8	67.2	79.4	51.7	181.6	118.8	85.7	96.4	91.1	173.1	176.5	174.8
Static-SiC (ours)	-	11.0	19.0	22.7	32.6	37.8	24.6	51.6	41.8	29.8	44.7	37.2	59.8	67.2	63.5
Dynamic-SiC (ours)	-	11.0	18.9	22.6	32.2	37.4	24.4	51.8	42.5	29.5	44.0	36.8	59.2	66.5	62.9

We denote the SiC-based on Static-/Dynamic-TUP is referred to as Static-/Dynamic-SiC, respectively. The proposed TUP facilitates the modeling of distribution differences among tasks and datasets, so the model focuses more on learning unified representations shared across tasks.

4.2. Model Architecture

As shown in Fig. 3, we build a skeleton bank, a collection of publicly available datasets including H3.6M, AMASS and 3DPW, covering common skeleton analysis tasks, containing input-target pairs formed as sequences of equal length, where prompts are randomly selected. Given a prompt pair $G_i^k = \{I_i^k, T_i^k\}$ and a query sample Q_j^k , we first map them to the C-dimensional feature space, and then splice I and T , Q and U respectively. The former represents our proposed Task-Guided Prompt, and the latter is the objective sequence of our focused analysis. The process can be formulated as:

$$\tilde{G} = \mathcal{E}(G_i^k) \in \mathbb{R}^{2F \times J \times C}, \quad (4)$$

$$\tilde{Q} = (\mathcal{E}(Q_j^k) \parallel U) \in \mathbb{R}^{2F \times J \times C}, \quad (5)$$

where $\mathcal{E}(\cdot)$ is a linear transformation and \parallel is concatenating operator. In essence, U represents the prior knowledge of the target skeleton sequence and is also an adapter that adapts to different tasks and different distributions of data for learning.

Drawing inspirations from [77], we connect spatial attention blocks and temporal attention blocks together in an alternating combination manner to form a two-stream

transformer $\mathcal{M}(\cdot)$ dedicated to capturing the dynamic characteristics of motion, which is implemented as:

$$\mathcal{M}(X) = \alpha^i \mathcal{T}_1^i \left(\mathcal{S}_1^i \left(X^{i-1} \right) \right) + \beta^i \mathcal{S}_2^i \left(\mathcal{T}_2^i \left(X^{i-1} \right) \right), \quad (6)$$

where α and β are balanced parameters determined adaptively. \mathcal{S}_1^i and \mathcal{T}_1^i are implemented using spatial self-attention mechanism and temporal self-attention mechanism, respectively. Additionally, $i = 1, \dots, N$.

Then, $\mathcal{M}(\cdot)$ takes \tilde{G} and \tilde{Q} as input and parallel iterates over them n_1 times respectively. Then, \tilde{G}^{n_1} and \tilde{Q}^{n_1} are integrated via aggregation function $\mathcal{A}(\cdot)$:

$$W^{n_1} = \mathcal{A}(\tilde{G}^{n_1}, \tilde{Q}^{n_1}), \quad (7)$$

we consider $\mathcal{A}(\cdot)$ to be a simple weighted summation function. The final output is obtained after iterating W n_2 times in $\mathcal{M}(\cdot)$, where $n_1 + n_2 = N$. Note that we take the second half of the output as the prediction.

4.3. Training and Inference

Training. We train our model in our proposed skeleton-centric in-context dataset mentioned in Sec. 3. In this training way, our Skeleton-in-Context can perform multiple tasks after one training. Since our output format is skeleton-like, following previous work [77], we adopt the 3D joint-level reconstruction loss:

$$Loss = \sum_{f=1}^F \sum_{j=1}^J \|P_{f,j} - G_{f,j}\|_2, \quad (8)$$

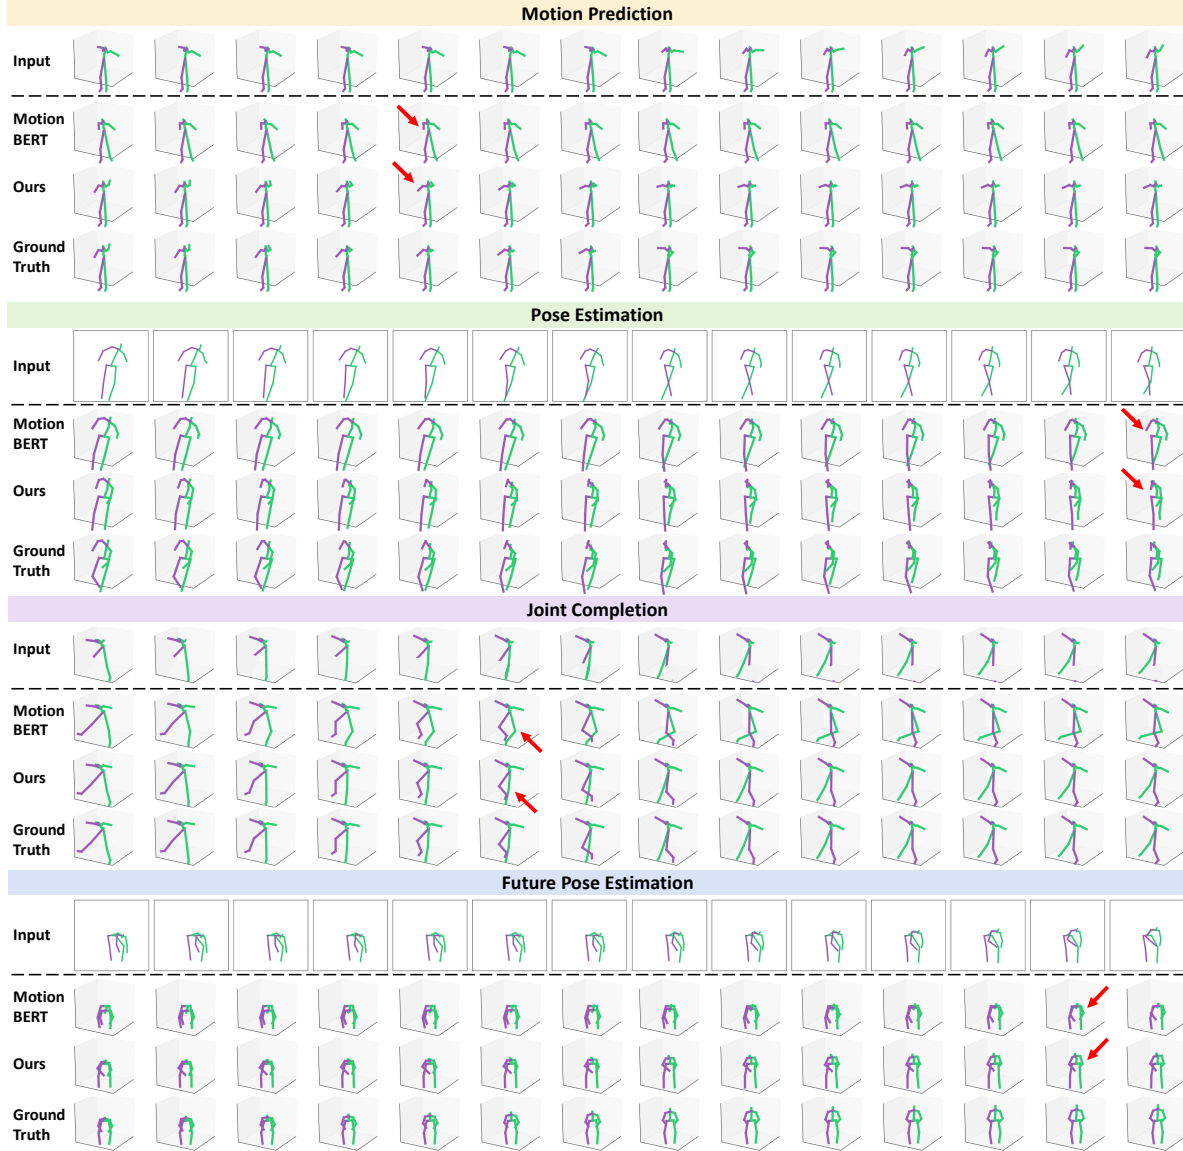


Figure 4. The comparison of visual results between our Skeleton-in-Context and recent SoTA method MotionBERT [77] on four tasks. Our method generates more accurate poses than MotionBERT, as can be seen from the spots with red arrows (↘) pointed to them.

Inference. After a training phase involving all tasks, we get an in-context model that focuses on dynamic skeleton sequences. It can perform multiple tasks at the same time according to the provided task-dependent prompt without any further parameter updates.

5. Experiments

Implementation Details. We implement the proposed Skeleton-in-Context model with the number of layers $N = 5$, number of attention heads $H = 8$, and hidden feature dimension $C = 256$. For each prompt input/target and query input/target, the sequence length is $F = 16$. All tasks are unified into a one-off, end-to-end training process without

any task-specific designs. The training takes about 6 hours for 120 epochs on 4 NVIDIA GeForce RTX 4090 GPUs.

Datasets and Metrics. H3.6M (Human3.6M) is used for the pose estimation and future pose estimation tasks. AMASS is used for the motion prediction task. 3DPW (3D Pose in the Wild) is used for the joint completion task. All of the tasks can be evaluated using Mean Per Joint Position Error.

5.1. Baseline Methods

In order to make a qualitative comparison with other methods in a relatively fair situation, we reimplement the task-specific SoTA methods in our proposed benchmark. Additionally, we transform recent methods into multi-task models and compare them with our model.

Table 2. Comparison results between task-specific models and our proposed models on generalization testing experiments on different datasets. We report Mean Per Joint Position Error (mm) for all tasks.

Methods	MP. (3DPW)						PE.		JC. (H3.6M)		
	40ms	80ms	120ms	160ms	200ms	Avg.	AMASS	3DPW	40%	60%	Avg.
Task-Specific Model: one architecture for one task											
EqMotion [65]	15.6	35.6	52.1	67.9	79.5	50.1	51.8	156.6	193.6	194.0	193.8
STCFormer [52]	18.3	39.3	58.9	75.0	88.6	56.0	45.0	141.2	177.2	178.4	177.8
siMLPe [41]	14.9	26.2	44.0	56.6	68.8	42.1	48.8	150.4	236.3	257.8	247.1
In-Context Model: one task-agnostic architecture for multiple tasks											
Static-SiC (ours)	7.9	21.5	37.7	54.0	69.1	38.1	22.0	76.4	162.5	171.9	167.2
Dynamic-SiC (ours)	7.1	21.3	37.4	53.5	68.0	37.5	21.7	76.0	150.1	151.7	151.0




Figure 5. Comparison of generalization. Our SiC completes the missing skeletons according to the customized prompt for an unseen task.

Copy Method. Following [5, 16], we copy the prompt target T_i^k as prediction results for comparison.

Multi-Task-Capable Models. For task-specific models [41, 52, 65] and multi-stage models [77], they are not able to handle multiple tasks at the same time. Therefore, we re-model them into end-to-end multi-task-capable models to achieve fair comparison.

In-Context Models. Since there are no in-context models excelling in processing multiple human perception tasks, we select the most relative in-context models with our benchmark, Point-In-Context [16] and PointMAE [43], which is proposed in 3D point cloud understanding. Specifically, we take each frame as a token in their framework. For PointMAE, we use the transformed version in [16].

5.2. Compared with SoTA Methods

Quantitative Results. We report experimental results on four tasks in Tab. 1. We term Skeleton-in-Context implemented with static/dynamic-TUP as Static/Dynamic-SiC, respectively. For other models, we reproduce them and report the results. As can be seen from Tab. 1, our proposed model outperforms all multi-task models, achieving state-of-the-art results on the in-context multi-task benchmark. Notably, the results of task-specific models on a single task reflect the difficulty of each task. However, we also exhibit impressive

performance compared to task-specific models, demonstrating that our proposed model is an expert in handling skeleton sequence-based multi-tasking.

Qualitative Results. We visualize and compare the results of our SiC and the most recent SoTA model, MotionBERT [77], which is re-structured as an end-to-end multi-task model for fair comparison. As highlighted in Fig. 4, our SiC can generate more accurate poses than MotionBERT according to the provided task-guided prompt.

5.3. Generalization Capability.

Generalize to New Datasets. To verify the generalization ability of our model, we conduct cross-dataset experiments on a specific task. For instance, our proposed model learns motion prediction on the AMASS dataset under the default benchmark, and then it is used to predict future motion sequences on 3DPW during testing. Similarly, we set up four experiments: motion prediction (AMASS→3DPW), pose estimation (H3.6M→AMASS, 3DPW), and joint completion (3DPW→H3.6M). As shown in Tab. 2, our model is able to generalize the learned task well to other datasets, while other task-specific models cannot accommodate the large bridge introduced across datasets even if they are specifically trained on this single task. Additionally, benefiting from the dynamic pseudo pose, the Dynamic-SiC exhibits

Table 3. Effectiveness of the Skeleton Prompt.

#	TGP	TUP	MP. MPJPE↓	PE. MPJPE↓	JC. MPJPE↓	FPE. MPJPE↓
1	✗	✗	111.7	159.2	543.7	160.2
2	✓	✗	25.0	52.5	37.7	65.2
3	✗	✓	26.0	52.4	37.8	78.4
4	✓	✓	24.4	51.8	36.8	62.9

Table 4. Different Prompting (TGP).

#	TGP	MP. MPJPE↓	PE. MPJPE↓	JC. MPJPE↓	FPE. MPJPE↓
1	Feature Sim.	25.2	52.1	37.7	64.1
2	Motion Sim.	25.2	52.1	37.8	64.1
3	Random	24.4	51.8	36.8	62.9

more excellent generalization ability than the Static-SiC. **Generalize to New Task.** Furthermore, We construct a prompt that has not appeared in the training set, which performs motion in-between. As Fig. 5 shows, our trained model can complete the missing skeletons according to the customized task-guided prompt. Our SiC generalizes well to a new task, which is unexplored in other works.

5.4. Ablation Study

Effectiveness of the Skeleton Prompt. We conduct ablation experiments on our proposed skeleton prompts. As shown in Tab. 3, the skeleton prompts solve the problem of directly extending previous work to skeleton sequences, which is falling into over-fitting. Meanwhile, applying both TGP and TUP to the model can achieve the best results. Our proposed skeleton prompts replace the random masking training strategy used in previous work. In particular, TUP creates a training environment that does not require the query target, which ensures that the model will not see the ground truth of the query input in advance.

Prompt Engineering. Since previous work has proven that different prompts in context learning affect the performance of the model, we conduct ablation experiments on the choice of prompts. As Tab. 4 shows, we select the prompt by the feature similarity and MPJPE minimum between the input skeleton sequence and the prompt, respectively. Experimental results show that the random selection method achieves the best results. Unlike 2D images and 3D point clouds, skeleton-based task data is highly similar, resulting in different prompts playing similar roles in the same task. While the random selection method expands the model’s receptive field to the data and achieves better results.

Architecture. We conduct ablation experiments on the architecture of the model. As shown in Tab. 5, after exploring different hidden feature dimensions and transformer depths, we find that the model performs best when the feature dimension is 256, and the depth is 5.

5.5. In-Context Initial Pose

We conduct an in-depth exploration of Dynamic-TUP. We extract the trained Dynamic-TUP weights, then randomly initialize a $(F, J, 3)$ tensor and use the trained encoder to in-

Table 5. Ablation on model architecture.

Arch.	Num.	MP. MPJPE↓	PE. MPJPE↓	JC. MPJPE↓	FPE. MPJPE↓
Dim.	64	29.3	57.1	45.8	80.5
	128	26.4	53.8	39.2	66.1
	256	24.4	51.8	36.8	62.9
	512	24.7	50.7	37.6	63.0
Dep.	3	25.1	53.7	39.2	65.6
	4	25.1	52.2	38.6	64.7
	5	24.4	51.8	36.8	62.9
	6	24.7	52.9	38.1	63.7




Figure 6. (a) The visualization of Static-TUP. (b) The concrete visual skeleton sequence of Dynamic-TUP.

crease its feature dimension to C-dimension. Finally, an optimizer is used to optimize the tensor and minimize its feature similarity with the trained Dynamic-TUP. The top of Fig. 6 represents the visualization of Static-TUP, while the bottom is the Dynamic-TUP. The concrete skeleton of dynamic-TUP is similar to the mean pose Static-TUP of the training set, which demonstrates that the Dynamic-TUP learned by our model is actually the prior knowledge whose receptive field is all ground truth in the training set. This prior knowledge helps our model learn more smoothly on multiple tasks, and adding learnable modules helps our model generalize to other datasets or new tasks (Sec. 5.3).

6. Conclusion

In this work, we propose the Skeleton-in-Context, designed to process multiple skeleton-base tasks simultaneously after just one training time. Specifically, we build a novel skeleton-based in-context benchmark covering various tasks. In particular, we propose skeleton prompts composed of TGP and TUP, which solve the overfitting problem of skeleton sequence data trained under the training framework commonly applied in previous 2D and 3D in-context models. Besides, we demonstrate that our model can generalize to different datasets and new tasks, such as motion completion. We hope our research builds the first step in the exploration of in-context learning for skeleton-based sequences, which paves the way for further research in this area.

Acknowledgement This work is supported by National Natural Science Foundation of China (No. 62203476), Natural Science Foundation of Shenzhen (No. JCYJ20230807120801002). This project is also supported by the interdisciplinary doctoral grants (iDoc 2021-360) from the Personalized Health and Related Technologies (PHRT) of the ETH domain, Switzerland.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 2022. **2**
- [2] Yutong Bai, Xinyang Geng, Kartikeya Mangalam, Amir Bar, Alan Yuille, Trevor Darrell, Jitendra Malik, and Alexei A Efros. Sequential modeling enables scalable learning for large vision models, 2023. **3**
- [3] Ivana Balažević, David Steiner, Nikhil Parthasarathy, Relja Arandjelović, and Olivier J Hénaff. Towards in-context scene understanding. *arXiv preprint arXiv:2306.01667*, 2023. **2, 3**
- [4] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. **2**
- [5] Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir Globerson, and Alexei Efros. Visual prompting via image inpainting. *NeurIPS*, 2022. **2, 3, 4, 5, 7**
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. **1, 2**
- [7] Haoye Cai, Chunyan Bai, Yu-Wing Tai, and Chi-Keung Tang. Deep video generation, prediction and completion of human action sequences. In *ECCV*, 2018. **2**
- [8] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *ICCV*, 2019. **2**
- [9] Yujun Cai, Yiwei Wang, Yiheng Zhu, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Chuanxia Zheng, Sijie Yan, Henghui Ding, et al. A unified 3d human motion synthesis model via conditional variational auto-encoder. In *ICCV*, 2021. **1, 3**
- [10] Yu Cheng, Bo Yang, Bo Wang, and Robby T Tan. 3d human pose estimation using spatio-temporal networks with explicit occlusion training. In *AAAI*, 2020. **2**
- [11] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *ICCV*, 2019. **2**
- [12] Yuanzheng Ci, Yizhou Wang, Meilin Chen, Shixiang Tang, Lei Bai, Feng Zhu, Rui Zhao, Fengwei Yu, Donglian Qi, and Wanli Ouyang. Unihcp: A unified model for human-centric perceptions. In *CVPR*, 2023. **2, 3**
- [13] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3d human motion prediction. In *CVPR*, 2020. **2**
- [14] Zhichao Deng, Xiangtai Li, Xia Li, Yunhai Tong, Shen Zhao, and Mengyuan Liu. Vg4d: Vision-language model goes 4d video recognition. In *ICRA*, 2024. **2**
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **2**
- [16] Zhongbin Fang, Xiangtai Li, Xia Li, Joachim M Buhmann, Chen Change Loy, and Mengyuan Liu. Explore in-context learning for 3d point cloud understanding. *arXiv preprint arXiv:2306.08659*, 2023. **1, 2, 3, 4, 5, 7**
- [17] Lin Geng Foo, Tianjiao Li, Hossein Rahmani, Qiuhong Ke, and Jun Liu. Unified pose sequence modeling. In *CVPR*, 2023. **1, 2, 3**
- [18] Jia Gong, Lin Geng Foo, Zhipeng Fan, Qiuhong Ke, Hossein Rahmani, and Jun Liu. Diffpose: Toward more reliable 3d pose estimation. In *CVPR*, 2023. **2, 3**
- [19] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to mlp: A simple baseline for human motion prediction. In *WACV*, 2023. **5**
- [20] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *TOG*, 2020. **2**
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. **2, 3, 4**
- [22] Gustav Eje Henter, Simon Alexanderson, and Jonas Beskow. Moglow: Probabilistic and controllable motion synthesis using normalising flows. *TOG*, 2020. **2**
- [23] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *ICCV*, 2019. **2**
- [24] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, 2013. **3, 5**
- [25] Mahdi M Kalayeh, Emrah Basaran, Muhittin Gökmen, Mustafa E Kamasak, and Mubarak Shah. Human semantic parsing for person re-identification. In *CVPR*, 2018. **3**
- [26] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. In *3DV*, 2020. **2**
- [27] Sameh Khamis, Cheng-Hao Kuo, Vivek K Singh, Vinay D Shet, and Larry S Davis. Joint learning for attribute-consistent person re-identification. In *ECCV*, 2015. **3**
- [28] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *CVPR*, 2018. **2**
- [29] Xiangtai Li, Henghui Ding, Wenwei Zhang, Haobo Yuan, Guangliang Cheng, Pang Jiangmiao, Kai Chen, Ziwei Liu, and Chen Change Loy. Transformer-based visual segmentation: A survey. *arXiv pre-print*, 2023. **2**
- [30] Xiaodan Liang, Ke Gong, Xiaohui Shen, and Liang Lin. Look into person: Joint body parsing & pose estimation network and a new benchmark. *PAMI*, 2018. **3**
- [31] Chunze Lin, Jiwen Lu, and Jie Zhou. Multi-grained deep feature learning for pedestrian detection. In *ICME*, 2018. **3**
- [32] Jun Liu, Henghui Ding, Amir Shahroudy, Ling-Yu Duan, Xudong Jiang, Gang Wang, and Alex C Kot. Feature boosting network for 3d pose estimation. *PAMI*, 2019. **2, 3**
- [33] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *PAMI*, 2019. **3**

- [34] Kenkun Liu, Rongqi Ding, Zhiming Zou, Le Wang, and Wei Tang. A comprehensive study of weight sharing in graph networks for 3d human pose estimation. In *ECCV*, 2020. 2, 3
- [35] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *CVPR*, 2018. 3
- [36] Jianxin Ma, Shuai Bai, and Chang Zhou. Pretrained diffusion models for unified human motion synthesis. *arXiv preprint arXiv:2212.02837*, 2022. 1, 3
- [37] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *ICCV*, 2019. 3
- [38] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *ICCV*, 2019. 2, 3
- [39] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *ECCV*, 2020. 2
- [40] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. 2
- [41] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017. 2, 5, 7
- [42] Xuecheng Nie, Jiashi Feng, and Shuicheng Yan. Mutual learning to adapt for joint human parsing and pose estimation. In *ECCV*, 2018. 3
- [43] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *ECCV*, 2022. 3, 4, 5, 7
- [44] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *CVPR*, 2019. 2
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3
- [46] Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. *arXiv preprint arXiv:2112.08633*, 2021. 2
- [47] Justyna Sarzynska-Wawer, Aleksander Wawer, Aleksandra Pawlak, Julia Szymanowska, Izabela Stefaniak, Michal Jarkiewicz, and Lukasz Okruszek. Detecting formal thought disorder by deep contextualized word representations. *Psychiatry Research*, 2021. 2
- [48] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016. 3
- [49] Wenkang Shan, Zhenhua Liu, Xinfeng Zhang, Shanshe Wang, Siwei Ma, and Wen Gao. P-stmo: Pre-trained spatial temporal many-to-one model for 3d human pose estimation. In *ECCV*, 2022. 2
- [50] Chi Su, Fan Yang, Shiliang Zhang, Qi Tian, Larry Steven Davis, and Wen Gao. Multi-task learning with low rank attribute embedding for multi-camera person re-identification. *PAMI*, 2017. 3
- [51] Yanpeng Sun, Qiang Chen, Jian Wang, Jingdong Wang, and Zechao Li. Exploring effective factors for improving visual in-context learning. *arXiv preprint arXiv:2304.04748*, 2023. 2, 3
- [52] Zhenhua Tang, Zhaofan Qiu, Yanbin Hao, Richang Hong, and Ting Yao. 3d human pose estimation with spatio-temporal criss-cross attention. In *CVPR*, 2023. 5, 7
- [53] Yonglong Tian, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Pedestrian detection aided by deep learning semantic tasks. In *CVPR*, 2015. 3
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3
- [55] Timo Von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018. 3
- [56] Chaoyang Wang, Xiangtai Li, Henghui Ding, Lu Qi, Jiangning Zhang, Yunhai Tong, Chen Change Loy, and Shuicheng Yan. Explore in-context segmentation via latent diffusion models. *arXiv*, 2024. 2
- [57] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3d pose estimation from videos. In *ECCV*, 2020. 2
- [58] Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images: A generalist painter for in-context visual learning. In *CVPR*, 2023. 1, 2, 3, 4
- [59] Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt: Towards segmenting everything in context. In *ICCV*, 2023. 2, 3
- [60] Xinshun Wang, Qiongjie Cui, Chen Chen, and Mengyuan Liu. Gcnxt: Towards the unity of graph convolutions for human motion prediction. In *AAAI*, 2024. 2
- [61] Xinshun Wang, Wanying Zhang, Can Wang, Yuan Gao, and Mengyuan Liu. Dynamic dense graph convolutional network for skeleton-based human motion prediction. *TIP*, 2024. 2
- [62] Zhendong Wang, Yifan Jiang, Yadong Lu, Yelong Shen, Pengcheng He, Weizhu Chen, Zhangyang Wang, and Mingyuan Zhou. In-context learning unlocked for diffusion models. *arXiv preprint arXiv:2305.01115*, 2023. 2
- [63] Jianzong Wu, Xiangtai Li, Shilin Xu, Haobo Yuan, Henghui Ding, Yibo Yang, Xia Li, Jiangning Zhang, Yunhai Tong, Xudong Jiang, Bernard Ghanem, and Dacheng Tao. Towards open vocabulary learning: A survey. *arXiv pre-print*, 2023. 2
- [64] Xiaoyang Wu, Zhuotao Tian, Xin Wen, Bohao Peng, Xihui Liu, Kaicheng Yu, and Hengshuang Zhao. Towards large-scale 3d representation learning with multi-dataset point prompt training. *arXiv preprint arXiv:2308.09718*, 2023. 2, 3
- [65] Chenxin Xu, Robby T Tan, Yuhong Tan, Siheng Chen, Yu Guang Wang, Xinchao Wang, and Yanfeng Wang. Eqmotion: Equivariant multi-agent motion prediction with invariant interaction reasoning. In *CVPR*, 2023. 5, 7
- [66] Shilin Xu, Xiangtai Li, Jingbo Wang, Guangliang Cheng, Yunhai Tong, and Dacheng Tao. Fashionformer: A simple, effective and unified baseline for human fashion segmentation and recognition. In *ECCV*, 2022. 3
- [67] Tianhan Xu and Wataru Takano. Graph stacked hourglass networks for 3d human pose estimation. In *CVPR*, 2021. 2

- [68] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *ICCV*, 2019. [2](#)
- [69] Angela Yao, Juergen Gall, and Luc Van Gool. Coupled action recognition and pose estimation from multiple views. *IJCV*, 2012. [3](#)
- [70] Bruce XB Yu, Zhi Zhang, Yongxu Liu, Sheng-hua Zhong, Yan Liu, and Chang Wen Chen. Gla-gcn: Global-local adaptive graph convolutional network for 3d human pose estimation from monocular video. In *ICCV*, 2023. [5](#)
- [71] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Jun-song Yuan. Mixste: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *CVPR*, 2022. [2](#)
- [72] Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. What makes good examples for visual in-context learning? *arXiv preprint arXiv:2301.13670*, 2023. [2](#), [3](#)
- [73] Weixi Zhao, Weiqiang Wang, and Yunjie Tian. Graformer: Graph-oriented transformer for 3d pose estimation. In *CVPR*, 2022. [2](#)
- [74] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *ICCV*, 2021.
- [75] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *ICCV*, 2021. [2](#)
- [76] Yi Zhou, Jingwan Lu, Connelly Barnes, Jimei Yang, Sitao Xiang, et al. Generative tweening: Long-term inbetweening of 3d human motions. *arXiv preprint arXiv:2005.08891*, 2020. [2](#)
- [77] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *ICCV*, 2023. [1](#), [3](#), [5](#), [6](#), [7](#)