

MotionAGFormer: Enhancing 3D Human Pose Estimation with a Transformer-GCNFormer Network

Soroush Mehraban^{1,2}, Vida Adeli^{1,3}, Babak Taati^{1,2,3}

¹KITE Research Institute, ²Institute of Biomedical Engineering, University of Toronto,

³Department of Computer Science, University of Toronto

{soroush.mehraban, vida.adeli}@mail.utoronto.ca, babak.taati@uhn.ca

Abstract

Recent transformer-based approaches have demonstrated excellent performance in 3D human pose estimation. However, they have a holistic view and by encoding global relationships between all the joints, they do not capture the local dependencies precisely. In this paper, we present a novel Attention-GCNFormer (AGFormer) block that divides the number of channels by using two parallel transformer and GCNFormer streams. Our proposed GCNFormer module exploits the local relationship between adjacent joints, outputting a new representation that is complementary to the transformer output. By fusing these two representation in an adaptive way, AGFormer exhibits the ability to better learn the underlying 3D structure. By stacking multiple AGFormer blocks, we propose MotionAGFormer in four different variants, which can be chosen based on the speed-accuracy trade-off. We evaluate our model on two popular benchmark datasets: Human3.6M and MPI-INF-3DHP. MotionAGFormer-B achieves state-of-the-art results, with P1 errors of 38.4 mm and 16.2 mm, respectively. Remarkably, it uses a quarter of the parameters and is three times more computationally efficient than the previous leading model on Human3.6M dataset. Code and models are available at <https://github.com/TaatiTeam/MotionAGFormer>.

1. Introduction

Human pose estimation in 3D space is an active area of research with significant implications for numerous applications, from augmented [22] and virtual reality [27] to autonomous vehicles [2, 9, 44], human-computer interaction [28] and beyond. With this vast range of applications, the demand for more accurate and computationally efficient pose estimation models continues to grow. In most real-world scenarios, pose sequences are captured in 2D, primarily due to the prevalent use of standard RGB cam-




Figure 1. Comparisons of recent pose uplifting methods on Human3.6M [15] (lower is better). MACs/frame denotes multiply-accumulate operations per each output frame. The proposed MotionAGFormer presents different variants and attains superior results, while maintaining computational efficiency.

eras. Consequently, one of the pivotal challenges in the field has been to effectively lift these 2D sequences into a 3D space. Accurate 3D human pose estimation enables the extraction of rich spatio-temporal information about human movements, and a deeper understanding of activities and interactions. Recent 3D lifting models leverage the inherent spatial and temporal coherence of human movements to enhance the precision of 3D pose predictions [41, 52, 55]. Nonetheless, despite the considerable advancements, there are several significant challenges that require attention.

The Transformer architecture [42], originally designed for NLP tasks, has been adapted to various computer vision problems, including pose estimation. Its ability to capture long-range dependencies and its innate self-attention mechanism make it a promising candidate for this domain. However, a sole reliance on global attention mechanisms, as employed by standard Transformers, may not be optimal for

pose estimation tasks. Human motion is inherently structured with local spatial and temporal dependencies.

One primary concern is the modeling of skeleton relations over time. Existing methods predominantly rely either on transformer architectures [41,49,55] or graph-based models [8,45,46]. While transformers excel at capturing long-term dependencies, graph models excel at local dependencies. So, there is an opportunity for a unified architecture that integrates the global perspective of transformers with the local precision of graph models.

Additionally, the race for achieving SOTA accuracy has often led to the development of increasingly complex models with a large number of parameters. Such models, despite their good accuracy, often become impractical for real-world applications where computational efficiency and swift response times are pivotal. Moreover, a predominant approach in recent models has been the prediction of a single 3D pose only for the central frame from a sequence of frames. This method, while seemingly efficient, leads to computational redundancy as it requires the reprocessing of several overlapping sequences. As a result, we instead employ a streamlined inference strategy that optimally exploits sequential data. This approach minimizes redundancy by predicting the complete 3D sequence of the input at a single forward pass.

In this paper, we introduce the MotionAGFormer, a novel transformer-graph hybrid architecture tailored for 3D human pose estimation. At its core, the MotionAGFormer harnesses the power of transformers to capture global information while simultaneously employing Graph Convolutional Networks (GCNs) to integrate local spatial and temporal relationships. We use an adaptive fusion to aggregate features from the transformer and graph streams. By doing so, we ensure a balanced and comprehensive representation of human motion, leading to enhanced accuracy in the 3D pose estimation (See Figure 1).

In summary, the main contributions of our paper are:

- *Novel Design:* We propose the *MotionAGFormer* model, in which we introduce a new GCNFormer module that excels in representing local dependencies inherent in human pose sequences.
- *Efficiency and Flexibility:* **i)** Our MotionAGFormer stands out due to its lightweight nature and faster speed with fewer parameters compared to previous SOTA methods, without compromising on accuracy. **ii)** Recognizing diverse needs, we offer different variants of MotionAGFormer, granting users the flexibility to make a balanced choice between accuracy and speed based on their specific requirements.
- *SOTA Performance:* MotionAGFormer achieves SOTA on two challenging datasets, Human3.6M and MPI-INF-3DHP.

2. Related works

3D human pose estimation. Current approaches to tackle this problem can be understood from two perspectives. From one perspective, models can be categorized based on the input video, which can be either multi-view or monocular. Models that rely on multi-view inputs [7, 16, 35, 36, 50] necessitate the simultaneous use of multiple cameras from different angles, which can be less feasible in real-world scenarios. From another perspective, considering their methodology, these models can be categorized as either direct 3D estimation approaches or 2D-3D lifting approaches. Direct 3D estimation methods [32, 33, 40, 54] infer the joints in 3D coordinate from the video frames without any intermediate step. Inspired by the rapid development and availability of accurate 2D pose estimation models, more recently, 2D-3D lifting methods first use an off-the-shelf 2D pose detectors [4, 29, 39] then lift 2D coordinates to 3D space [41, 49, 52, 53, 55]. In this work, we are using 2D-3D lifting methods by having monocular video as the input.

Transformer-based methods. Transformers [42] have shown promising result in different visual tasks [10, 38, 47, 56]. In the field of 3D human pose estimation, PoseFormer [53] was the first purely transformer-based model. PoseFormerV2 [52] improved its computational efficacy by employing a frequency-domain representation that also made it robust against sudden movements in noisy data. MHFormer [21] addressed the problem of self-occlusion and depth ambiguity by proposing a module that learns multiple plausible pose hypotheses. P-STMO [37] proposed masked pose modeling and reduced the final error by self-supervised pretraining the model. Enfalt *et al.* [11] decreased the computational complexity by leveraging masked token modeling. StridedFormer [20] replaced fully-connected layers in the feed-forward network of the transformer encoder with strided convolutions to progressively shrink the sequence length and efficiently lift the central frame. Unlike the abovementioned models that estimate the 3D pose for only the center frame of a sequence, MixSTE [49] provided 3D estimates for every frame in the input sequence. STCFormer [41] decreased computational complexity by separating the correlation learning into spatial and temporal components. HSTFormer [34] proposed hierarchical transformer encoders for better capturing spatial-temporal correlations. In addition to joint-joint attention which is commonly used, HDFormer [3] included bone-joint and hyperbone-joint interactions. Some works exploited the output motion representation for various tasks. MotionBERT [55] fine-tuned the model learned for the task of 3D human pose estimation for tasks such as action recognition and 3D human mesh recovery, while UPS [12] trained a unified model for action recognition, 3D pose estimation, and early action prediction at the same time.

Graph Convolutional Network. GCN-based methods have achieved remarkable success within the domain of skeleton-based action recognition [5, 19, 23]. Despite their computational efficiency in 3D human pose estimation [6, 8, 43, 51], they usually cannot show competitive error compared to transformer-based counterparts. This is primarily due to their focus on local joints alone. Recently, GLA-GCN [46] introduced an adaptive GCN approach to leverage global representation. By employing a strided design to reduce its temporal scope, they achieve competitive 3D human pose estimation against various transformer-based models, all while maintaining a lighter memory load. However, the effectiveness of the proposed module in extracting global representation is not on par with that of an attention module.

Hybrid methods. These methods use different modules to capture distinct aspects of the input sequence and are not extensively explored yet. Recently, DC-GCT [17] and GFormer [13] proposed leveraging GCN modules in addition to transformers. However, since DC-GCT originally designed for a single frame input and GFormer showed sub-optimal performance by sequentially aggregating the modules, they fall short in delivering competitive outcomes when contrasted with transformer-based methods.

3. Method

3.1. Preliminary

We begin this section by reviewing the concept of MetaFormer [48], which forms the core of our encoders. A MetaFormer can be described as a generalization of the Transformer architecture [42], wherein the attention module is substituted with any mixer capable of transforming information among tokens. Specifically, for an input $X \in \mathbb{R}^{N \times C}$, with N denoting the token numbers and C representing the embedding dimension, the token mixer can be formally expressed as

$$Y = \text{TokenMixer}(\text{Norm}(X)) + X, \quad (1)$$

where $\text{Norm}(\cdot)$ denotes a normalization method such as batch or layer normalization [1, 14], and $\text{TokenMixer}(\cdot)$ denotes a module that combines information among tokens. Our approach uses two parallel token mixers: Multi-head Self-Attention (MHSA) and Graph Convolutional Networks (GCNs) [18], each contributing uniquely to the information transformation process.

3.2. Overall architecture

Our objective is to lift a 2D (pixel coordinate) skeleton sequence to accurate 3D pose sequences. To this end, we propose the MotionAGFormer architecture, which uses both attention (Transformer) and graph convolutional (GC-

NFormer) to lift motion sequences. An overview of this architecture is shown in Figure 2a.

The model takes a 2D input sequence with confidence score $\mathbf{X} \in \mathbb{R}^{T \times J \times 3}$, where T and J refer to the number of frames and joint numbers, respectively. It then proceeds to map each joint in each time frame to a d -dimensional feature, $\mathbf{F}^{(0)} \in \mathbb{R}^{T \times J \times d}$, using a linear projection layer. Then a spatial position embedding $\mathbf{P}_{\text{pos}}^s \in \mathbb{R}^{1 \times J \times d}$ is added to the tokens. It is important to highlight here that our model does not disregard the temporal token order as that information is preserved in the GCNFormer stream (further discussion in ablation studies 4.5).

Subsequent to position embedding, we use N blocks of AGFormer (Section 3.3) to compute $\mathbf{F}^{(i)} \in \mathbb{R}^{T \times J \times d}$ ($i = 1, \dots, N$) to effectively capture the underlying 3D structure of the skeleton sequence. Finally, we map $\mathbf{F}^{(N)}$ to a higher dimension by applying a linear layer and \tanh activation to compute motion semantic $\mathbf{M} \in \mathbb{R}^{T \times J \times d'}$ and use a regression head to estimate 3D pose $\hat{\mathbf{P}} \in \mathbb{R}^{T \times J \times 3}$. The lifting loss contains position (L_{3D}) and velocity ($L_{\Delta P}$) terms defined as

$$\begin{aligned} L_{3D} &= \sum_{t=1}^T \sum_{j=1}^J \|\hat{\mathbf{P}}_{t,j} - \mathbf{P}_{t,j}\|, \\ L_{\Delta P} &= \sum_{t=2}^T \sum_{j=1}^J \|\Delta \hat{\mathbf{P}}_{t,j} - \Delta \mathbf{P}_{t,j}\|, \end{aligned} \quad (2)$$

where $\Delta \hat{\mathbf{P}}_t = \hat{\mathbf{P}}_t - \hat{\mathbf{P}}_{t-1}$, $\Delta \mathbf{P}_t = \mathbf{P}_t - \mathbf{P}_{t-1}$. The total lifting loss is then defined as

$$L = L_{3D} + \lambda_{\Delta P} L_{\Delta P}, \quad (3)$$

where the constant coefficient $\lambda_{\Delta P}$ is used to balance position accuracy and motion smoothness.

3.3. AGFormer Block

The AGFormer block uses a dual-stream architecture. Each stream consists of two components: a Spatial MetaFormer (Figure 2b) followed by a Temporal MetaFormer (Figure 2c). The Spatial MetaFormer processes individual body joints as distinct tokens, effectively capturing intra-frame relationships within a single frame. The Temporal MetaFormer, on the other hand, treats each frame as a single token, thus capturing inter-frame relationships over time. The key distinction between the two streams lies in their token mixer type. While one stream employs Transformers, the other stream uses GCNFormers.

Transformer stream. This stream employs a Spatial Multi-Head Self-Attention (S-MHSA) to capture spatial relationships, followed by a Temporal Multi-Head Self-Attention (T-MHSA) to capture temporal relationships. The S-MHSA is defined as

$$\begin{aligned} \text{S-MHSA}(\mathbf{Q}_s, \mathbf{K}_s, \mathbf{V}_s) &= \text{Concat}(\text{head}_1, \dots, \text{head}_h) \mathbf{W}_s^{(O)}, \\ \text{head}_i &= \text{softmax}\left(\frac{\mathbf{Q}_s^{(i)} (\mathbf{K}_s^{(i)})^T}{\sqrt{d_k}}\right) \mathbf{V}_s^{(i)}, \end{aligned} \quad (4)$$




Figure 2. (a) **MotionAGFormer** is a novel architecture featuring N dual-stream spatio-temporal blocks, wherein one stream employs Transformers and the other leverages GCNFormers. (b) **Spatial MetaFormer**. Each input token represents an individual joint of the human body. (c) **Temporal MetaFormer**. Input tokens are frames of pose sequence.

where $\mathbf{W}_s^{(O)}$ is a projection parameter matrix, h is the number of parallel attention heads, and d_k is the feature dimension of \mathbf{K}_s . For computing the query matrix \mathbf{Q}_s , the key matrix \mathbf{K}_s , and the value matrix \mathbf{V}_s , we have

$$\mathbf{Q}_s^i = \mathbf{F}_s \mathbf{W}_s^{(Q,i)}, \mathbf{K}_s^i = \mathbf{F}_s \mathbf{W}_s^{(K,i)}, \mathbf{V}_s^i = \mathbf{F}_s \mathbf{W}_s^{(V,i)}, \quad (5)$$

where $\mathbf{F}_s \in \mathbb{R}^{B \times T \times J \times d}$ is spatial feature and $\mathbf{W}_s^{(Q,i)}$, $\mathbf{W}_s^{(K,i)}$, $\mathbf{W}_s^{(V,i)}$ are projection matrices and B is the batch size. The S-MHSA result is subsequently fed into a multilayer perceptron (MLP), followed by a residual connection and LayerNorm. This completes the first MetaFormer, i.e. the Spatial Transformer.

Next, we reshape \mathbf{F}_s into $\mathbf{F}_T \in \mathbb{R}^{B \times J \times T \times d}$ to prepare per-joint temporal feature as the input of T-MHSA. Here we have

$$\begin{aligned} \text{T-MHSA}(\mathbf{Q}_T, \mathbf{K}_T, \mathbf{V}_T) &= \text{Concat}(\text{head}_1, \dots, \text{head}_h) \mathbf{W}_T^{(O)}, \\ \text{head}_i &= \text{softmax}\left(\frac{\mathbf{Q}_T^{(i)} (\mathbf{K}_T^{(i)})^T}{\sqrt{d_k}}\right) \mathbf{V}_T^{(i)}, \end{aligned} \quad (6)$$

where \mathbf{Q}_T , \mathbf{K}_T , and \mathbf{V}_T are calculated similar to Eqn. (5).

GCNFormer stream. Unlike the Transformer stream, which aggregates global information, the GCNFormer stream focuses on local spatial and temporal relationships present within the skeleton sequence. While the local information is also available to the Transformers, the inclusion of this parallel stream allows the model to more effectively balance the integration of local and global information (see ablation analysis 4.5). The customized GCN module [25] used in our GCNFormer is defined as:

$$\text{GCN}(\mathbf{F}^{(i)}) = \sigma(\mathbf{F}^{(i)} + \text{Norm}(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} \mathbf{F}^{(i)} W_1 + \mathbf{F}^{(i)} W_2)). \quad (7)$$

Where $\tilde{A} = A + I_N$ represents the adjacency matrix with self-connections added, I_N stands for the identity matrix, $\tilde{D}ii = \sum_j \tilde{A}jj$ is defined as the sum of elements along the diagonal of \tilde{A} , and W_1 , W_2 denote trainable weight matrices specific to each layer. The activation function $\sigma(\cdot)$, such as ReLU, is applied, along with Batch Normalization [14]. The GCN’s output is then passed through an MLP, followed by residual connection and LayerNorm.

The difference between the Spatial GCNFormer and the Temporal GCNFormer lies in their adjacency matrices and input features. The input features resemble that of the Transformer stream. In the Spatial GCNFormer, the adjacency matrix represents the human topology (Figure 3a). For Temporal GCNFormer, on the other hand, we calculate the similarity between a single joint at different time frames using $\text{Sim}(\mathbf{F}_T^{t_i}, \mathbf{F}_T^{t_j}) = (\mathbf{F}_T^{t_i})^T \mathbf{F}_T^{t_j}$ and choose the K nearest neighbors as the connected nodes in the graph (Figure 3b). Hence, the graph topology in Temporal GCNFormer is determined by the learned node features.

Adaptive Fusion. Similar to MotionBERT [55], we use adaptive fusion to aggregate extracted features of the Transformer and GCNFormer streams. This is defined as:

$$\mathbf{F}^{(i)} = \alpha_{TF} \mathbf{F}^{(i)} \circ \mathbf{F}_{TF}^{(i-1)} + \alpha_{GF} \mathbf{F}^{(i)} \circ \mathbf{F}_{GF}^{(i-1)}, \quad (8)$$

where $\mathbf{F}^{(i)}$ represents the feature embedding extracted at depth i , the element-wise multiplication denoted by \circ , and $\mathbf{F}_{TF}^{(i-1)}$, $\mathbf{F}_{GF}^{(i-1)}$ refer to the extracted Transformer stream and GCNFormer stream features at depth $i - 1$, respectively. The adaptive fusion weights α_{TF} and α_{GF} are




Figure 3. GCNFormer module topology. (a) The Spatial GCNFormer employs the human skeleton as its underlying topology. (b) The Temporal GCNFormer uses K-nearest neighbor (K-NN) to determine connected edges by considering the highest similarity of each joint (e.g., left wrist and right ankle in the figure) across the entire time frame. After K-NN, each row is connected to K columns.

defined as

$$\alpha_{TF}^{(i)}, \alpha_{GF}^{(i)} = \text{softmax}(W \cdot \text{Concat}(\mathbf{F}_{TF}^{(i-1)}, \mathbf{F}_{GF}^{(i-1)})), \quad (9)$$

where W is a learnable linear transformation.

4. Experiments

We evaluate the performance of our proposed MotionAGFormer on two large-scale 3D human pose estimation datasets, i.e., Human3.6M [15] and MPI-INF-3DHP [26].

4.1. Datasets and Evaluation Metrics

Human3.6M is a widely used indoor dataset for 3D human pose estimation. It contains 3.6 million video frames of 11 subjects performing 15 different daily activities. To ensure fair evaluation, we follow the standard approach and train the model using data from subjects 1, 5, 6, 7, and 8, and then test it on data from subjects 9 and 11. Following previous works [41, 52], we use two protocols for evaluation. The first protocol (referred to as P1) uses Mean Per Joint Position Error (MPJPE) in millimeters between the estimated pose and the actual pose, after aligning their root joints (sacrum). The second protocol (referred to as P2) measures Procrustes-MPJPE, where the actual pose and the estimated pose are aligned through a rigid transformation.

MPI-INF-3DHP is another large-scale dataset gathered in three different settings: green screen, non-green screen, and outdoor environments. Following previous works [41, 52], MPJPE, Percentage of Correct Keypoint (PC) within 150 mm range, and Area Under the Curve (AUC) are reported as evaluation metrics.

4.2. Implementation Details

Model Variants. We build four different configurations of our model, as summarized in Table 1. Our base model, known as MotionAGFormer-B, strikes a balance between accurate estimation and computational cost. The remaining variants are named according to their parameter size and computational demands and selection of each variant can be based on an application’s requirements, such as choosing between real-time processing or more precise estimations. The motion semantic dimension is $d' = 512$, the expansion layer of each MLP is $\alpha = 4$, the number of attention heads is $h = 8$, and the number of temporal neighbours in GCNformer stream is $k = 2$, for all experiments.

Table 1. Details of MotionAGFormer model variants. N : Number of layers. d : Hidden size. T : Number of input frames.

Method	N	d	T	Params	MACs
MotionAGFormer-XS	12	64	27	2.2 M	1.0 G
MotionAGFormer-S	26	64	81	4.8 M	6.6 G
MotionAGFormer-B	16	128	243	11.7 M	48.3 G
MotionAGFormer-L	26	128	243	19.0 M	78.3 G

Experimental settings. Our model is implemented using PyTorch [31] and executed on a setup with two NVIDIA A40 GPUs. We apply horizontal flipping augmentation for both training and testing following [52, 55]. For model training, we set each mini-batch as 16 sequences. The network parameters are optimized using AdamW [24] optimizer over 90 epochs with a weight decay of 0.01. The initial learning rate is set to $5e-4$ with an exponential learning rate decay schedule and the decay factor is 0.99. We use the Stacked Hourglass [30] 2D pose detection results and 2D ground truths on Human3.6M, following [55]. For MPI-INF-3DHP, ground truth 2D detection is used following a similar approach used in comparison baselines [41, 52].

4.3. Performance comparison on Human3.6M

We compare our MotionAGFormer with other models on Human3.6M. For a fair comparison, only results of models without extra pre-training on additional data is included. The results, outlined in Table 2, demonstrate that MotionAGFormer-L attains a P1 error of 38.4 mm for estimated 2D pose and 17.3 mm for ground truth 2D pose. Remarkably, this is achieved with approximately half the computational requirements and parameters compared to the previous SOTA (MotionBERT [55]), while being 0.8 mm and 0.5 mm more accurate, respectively. When comparing the Base and Large variants, MotionAGFormer-L shares the same P1 error as MotionAGFormer-B in the presence of noisy data, while exhibiting a 2.1 mm reduction in error when using ground truth 2D data. Furthermore,




Figure 4. Qualitative comparisons of MotionAGFormer with STCformer [41], PoseFormerV2 [52], and MotionBERT [55]. The transparent gray skeleton is the ground-truth 3D pose. Blue represents the torso and left part, and Red indicates the right part of the estimated body.

Table 2. Quantitative comparisons on Human3.6M. T : Number of input frames. CE: Estimating center frame only. P1: MPJPE error (mm). P2: P-MPJPE error (mm). $P1^\dagger$: P1 error on 2D ground truth. (*) denotes using HRNet [39] for 2D pose estimation. The best and second-best scores are in bold and underlined, respectively. For per action result, refer to the supplementary material.

Method	T	CE	Param	MACs	MACs/frame	$P1\downarrow$ / $P2\downarrow$	$P1^\dagger\downarrow$
*MHFormer [21] CVPR'22	351	✓	30.9 M	7.0 G	7,096 M	43.0/34.4	30.5
MixSTE [49] CVPR'22	243	×	33.6 M	139.0 G	572 M	40.9/32.6	21.6
P-STMO [37] ECCV'22	243	✓	6.2 M	0.7 G	740 M	42.8/34.4	29.3
Stridedformer [20] TMM'22	351	✓	4.0 M	0.8 G	801 M	43.7/35.2	28.5
Einfalt <i>et al.</i> [11] WACV'23	351	✓	10.4 M	0.5 G	498 M	44.2/35.7	-
STCFormer [41] CVPR'23	243	×	4.7 M	19.6 G	80 M	41.0/ <u>32.0</u>	21.3
STCFormer-L [41] CVPR'23	243	×	18.9 M	78.2 G	321 M	40.5/ 31.8	-
PoseFormerV2 [52] CVPR'23	243	✓	14.3 M	0.5 G	528 M	45.2/35.6	-
UPS [12] CVPR'23	243	✓	-	-	-	40.8/32.5	-
GLA-GCN [46] ICCV'23	243	✓	1.3 M	1.5 G	1,556 M	44.4/34.8	21.0
MotionBERT [55] ICCV'23	243	×	42.5 M	174.7 G	719 M	<u>39.2</u> / <u>32.9</u>	<u>17.8</u>
HDFormer [3] IJCAI'23	96	×	3.7 M	0.6 G	6 M	42.6/33.1	21.6
HSTFormer [34] arXiv'23	81	×	22.7 M	1.0 G	13 M	42.7/33.7	27.8
DC-GCT [17] arXiv'23	81	✓	3.1 M	41 M	41 M	44.7/-	-
MotionAGFormer-XS	27	×	2.2 M	1.0 G	37 M	45.1/36.9	28.1
MotionAGFormer-S	81	×	4.8 M	6.6 G	81 M	42.5/35.3	26.5
MotionAGFormer-B	243	×	11.7 M	48.3 G	198 M	38.4 /32.6	19.4
MotionAGFormer-L	243	×	19.0 M	78.3 G	322 M	38.4 / <u>32.5</u>	17.3

MotionAGFormer-S takes in only a third of the frames compared to the baselines, yet it manages to attain a superior P1 error compared to a majority of them. Similarly, MotionAGFormer-XS delivers comparable performance to PoseTransformerV2, even though it receives only a ninth of the input information (27 vs. 243 frames) and operates with approximately seven times fewer parameters.

4.4. Performance comparison on MPI-INF-3DHP

In evaluating our method on the MPI-INF-3DHP dataset, we modified our base and large variants to use 81 frames due to shorter video sequences. Across all variants, our method consistently outperforms others in terms MPJPE. Notably, our large variant achieves remarkable results with

an 85.3% AUC and a 16.2 mm P1 error. This outperforms the best models by a significant margin of 1.4% in AUC and 6.9 mm in P1 error. However, it achieves 98.2% PCK, which is 0.5% lower than the PCK performance of the compared models (Table 3).

Table 3. Quantitative comparisons on MPI-INF-3DHP. T : Number of input frames. The best and second-best scores are in bold and underlined, respectively.

Method	T	PCK \uparrow	AUC \uparrow	MPJPE \downarrow
MHFormer [21]	9	93.8	63.3	58.0
MixSTE [49]	27	94.4	66.5	54.9
P-STMO [37]	81	97.9	75.8	32.2
Einfalt <i>et al.</i> [11]	81	95.4	67.6	46.9
STCFormer [41]	81	98.7	83.9	23.1
PoseFormerV2 [52]	81	97.9	78.8	27.8
GLA-GCN [46]	81	<u>98.5</u>	79.1	27.7
HSTFormer [34]	81	97.3	71.5	41.4
HDFormer [3]	96	98.7	72.9	37.2
MotionAGFormer-XS	27	98.2	83.5	19.2
MotionAGFormer-S	81	98.3	<u>84.5</u>	<u>17.1</u>
MotionAGFormer-B	81	98.3	84.2	18.2
MotionAGFormer-L	81	98.2	85.3	16.2

4.5. Ablation Studies

A series of ablation studies were conducted on the Human3.6M dataset to explore different design choices for the AGFormer block.

The initial ablation study investigates the influence of favoring the number of AGFormer blocks and the width of our model on the P1 error. As shown in Table 4, the trend generally leans towards favoring a model that is deeper but narrower. Interestingly, a model with 16 AGFormer blocks and a width of 128 exhibits similar performance to a model featuring 12 AGFormer blocks with a width of 256, while the first configuration uses approximately three times less memory and computational resources.

The second part of the ablation study explores alternative modules for the graph stream within the AGFormer block. The results are presented in Table 5. Shifting from GCNFormer to GCN blocks degrades the P1 error by 0.7 mm. Similarly, substituting the Temporal GCNFormer with TCN degrades the P1 error by 0.2 mm when using GCN and 0.9 mm when using GCNFormer. Lastly, replacing the GCN with a CTR-GCN [5] increases the P1 error by 2.5 mm. We hypothesize that is because of the continuous refinement of topology for each channel in CTR-GCN during training. As a result, every joint not only obtains information from its immediate neighbors but also from all other joints. This prevents the generation of complementary information for the transformer stream.

To explore the impact of positional embedding on the

the final performance, we conducted a series of experiments outlined in Table 6. Surprisingly, including temporal positional embedding in addition to spatial positional embedding leads to a 0.6 mm increase in P1 error. As mentioned in Section 3.2, this result stems from the non-permutation equivariant nature of the GCNFormer stream. Unlike transformers, our network inherently maintains the temporal sequence of frames. However, as shown in Figure 5, adding temporal positional embedding leads to a better acceleration error (0.88 mm) compared to using spatial positional embedding (0.94 mm).

Finally, to verify the efficiency of the proposed AGFormer block, Table 7 shows alternative blocks. When using GCNFormer, a P1 error of 57.5 mm is observed, indicating its limited capability to accurately capture the underlying 3D sequence structure. Nevertheless, a hybrid approach involving both GCNFormer and Transformer yields a noteworthy improvement, reducing the P1 error by 5.2 mm compared to using Transformer alone. Furthermore, the sequential fusion of these two modules is not as effective as their parallel integration.

Table 4. The P1 error comparison by varying number of AGFormer blocks and number of channels. d : Number of channels in each AGFormer block. d' : Number of channels before regression head. T is kept 243 in all experiments.

Layers	d	d'	Param	MACs	P1
5	512	512	58 M	252.8 G	39.2
4	256	512	11.7 M	54.6 G	40.2
6	256	512	17.5 M	81.64 G	39.6
8	256	512	23.3 M	108.6 G	39.5
10	256	512	29.1 M	135.7 G	38.6
12	256	512	34.9 M	162.7 G	38.4
14	256	512	40.7 M	189.7 G	38.6
16	128	512	11.7 M	64.7 G	38.4
16	128	256	11.7 M	64.5 G	38.7
16	128	1024	11.7 M	64.9 G	39.2
26	64	512	4.8 M	39.5 G	39.5

Table 5. The P1 error comparison when using alternative modules in graph stream of AGFormer block.

GCN	TCN	GCNFormer	CTR-GCNFormer	P1
✓	-	-	-	39.1
✓	✓	-	-	39.3
-	✓	✓	-	39.7
-	-	-	✓	40.9
-	-	✓	-	38.4

4.6. Qualitative Analysis

Validation of MotionAGFormer-B is conducted using the adjacency matrix of temporal GCNFormer and visualization of 3D human pose estimation. The instances for vali-




Figure 5. **Acceleration error when using different type of positional embedding.** Using temporal positional embedding leads to 0.88 mm error and a faster convergence compared to spatial positional embedding, without embedding, and both embeddings that reach 0.94 mm, 0.90 mm, and 0.97 mm respectively.

data are randomly chosen from the evaluation set of Human3.6M.

Qualitative Comparisons. Figure 4 compares MotionAGFormer-B with recent approaches including STCFormer [41], PoseFormerV2 [52], and MotionBERT [55]. By design, confidence scores of the 2D detector are only included into MotionBERT and MotionAGFormer-B. Overall, MotionAGFormer-B shows better reconstruction results across all three samples than PoseFormerV2 and STCFormer, while maintaining competitive performance with MotionBERT. Specifically, in Figure 4a, MotionAGFormer exhibits improved alignment with the ground truth in comparison to alternative approaches. In Figure 4b, it displays a slightly superior alignment, whereas in Figure 4c, its alignment is marginally less optimal than that of MotionBERT.

Temporal adjacency visualization. The adjacency matrix of temporal GCNFormer is visualized in Figure 6. In the initial layers, individual joints exhibit distinct adjacency matrices, which should vary across different sequences due to the changing joint positions over time. Nevertheless, as we progress through the model’s depth, it appears to learn a representation where each joint is most akin to its own state in neighboring frames. Consequently, this leads to connections being formed with adjacent frames.

5. Conclusion

We introduced MotionAGFormer, a novel approach that leverages GCNFormer to capture intricate local joint relationships, and combines it with Transformer that effectively




Figure 6. Temporal adjacency matrix of a random sequence on three different joints, from the Human3.6M dataset, at the first layer (left), middle layer (center), and last layer (right). K in K-NN is set to 2.

Table 6. The P1 error comparison when using different positional embedding.

Temporal Embedding	Spatial Embedding	P1
-	-	39.3
-	✓	38.4
✓	-	38.9
✓	✓	40.5

Table 7. **Comparison of different MetaFormer integration.** All the models are trained on Human3.6M with our MotionAGFormer-B settings.

Method	P1
GCNFormer only	57.5
Transformer only	43.6
GCNFormer → Transformer (Sequential)	39.1
Transformer → GCNFormer (Sequential)	38.9
Transformer - GCNFormer (Parallel)	38.4

captures global joint interdependencies. This fusion enhances the model’s ability to comprehend the inherent 3D structure within input 2D sequences. Additionally, MotionAGFormer offers various adaptable variants, allowing the selection of an optimal balance between speed and accuracy. Empirical evaluations show that our method surpasses alternative methods on Human3.6M and MPI-INF-3DHP.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. **3**
- [2] Peter Bauer, Arij Bouazizi, Ulrich Kressel, and Fabian B Flohr. Weakly supervised multi-modal 3d human body pose estimation for autonomous driving. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–7. IEEE, 2023. **1**
- [3] Hanyuan Chen, Jun-Yan He, Wangmeng Xiang, Wei Liu, Zhi-Qi Cheng, Hanbing Liu, Bin Luo, Yifeng Geng, and Xuansong Xie. HDFormer: High-order directed transformer for 3d human pose estimation. *arXiv preprint arXiv:2302.01825*, 2023. **2, 6, 7**
- [4] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7103–7112, 2018. **2**
- [5] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13359–13368, 2021. **3, 7**
- [6] Hongsuk Choi, Gyeongseok Moon, and Kyoung Mu Lee. Pose2Mesh: Graph convolutional network for 3d human pose and mesh recovery from a 2d human pose. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 769–787. Springer, 2020. **3**
- [7] Sungho Chun, Sungbum Park, and Ju Yong Chang. Learnable human mesh triangulation for 3d human pose and shape estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2850–2859, 2023. **2**
- [8] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing network structure for 3d human pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2262–2271, 2019. **2, 3**
- [9] Phillip Czech, Markus Braun, Ulrich Kreßel, and Bin Yang. On-board pedestrian trajectory prediction using behavioral features. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 437–443. IEEE, 2022. **1**
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **2**
- [11] Moritz Einfalt, Katja Ludwig, and Rainer Lienhart. Uplift and upsample: Efficient 3d human pose estimation with uplifting transformers. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2023. **2, 6, 7**
- [12] Lin Geng Foo, Tianjiao Li, Hossein Rahmani, Qiuhong Ke, and Jun Liu. Unified pose sequence modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13019–13030, 2023. **2, 6**
- [13] Rongtian Huo, Qing Gao, Jing Qi, and Zhaojie Ju. 3d human pose estimation in video for human-computer/robot interaction. In *International Conference on Intelligent Robotics and Applications*, pages 176–187. Springer, 2023. **3**
- [14] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. **3, 4**
- [15] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. **1, 5**
- [16] Karim Isakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. Learnable triangulation of human pose. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7718–7727, 2019. **2**
- [17] Hongbo Kang, Yong Wang, Mengyuan Liu, Doudou Wu, Peng Liu, and Wenming Yang. Double-chain constraints for 3d human pose estimation in images and videos. *arXiv preprint arXiv:2308.05298*, 2023. **3, 6**
- [18] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. **3**
- [19] Jungho Lee, Minhyeok Lee, Dogyoon Lee, and Sangyoun Lee. Hierarchically decomposed graph convolutional networks for skeleton-based action recognition. *arXiv preprint arXiv:2208.10741*, 2022. **3**
- [20] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, Pichao Wang, and Wenming Yang. Exploiting temporal contexts with strided transformer for 3d human pose estimation. *IEEE Transactions on Multimedia*, 25:1282–1293, 2022. **2, 6**
- [21] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. MHFormer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13147–13156, 2022. **2, 6, 7**
- [22] Hwei-Yung Lin and Ting-Wen Chen. Augmented reality with human body interaction based on monocular 3d pose estimation. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 321–331. Springer, 2010. **1**
- [23] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 143–152, 2020. **3**
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. **5**
- [25] Cheng Luo, Siyang Song, Weicheng Xie, Linlin Shen, and Hatice Gunes. Learning multi-dimensional edge feature-based au relation graph for facial action unit recognition. *arXiv preprint arXiv:2205.01782*, 2022. **4**
- [26] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild

- using improved cnn supervision. In *2017 international conference on 3D vision (3DV)*, pages 506–516. IEEE, 2017. [5](#)
- [27] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. VNet: Real-time 3d human pose estimation with a single rgb camera. *Acm transactions on graphics (tog)*, 36(4):1–14, 2017. [1](#)
- [28] Tewodros Legesse Munea, Yalew Zelalem Jembre, Halefom Tekle Weldegebriel, Longbiao Chen, Chenxi Huang, and Chenhui Yang. The progress of human pose estimation: A survey and taxonomy of models applied in 2d human pose estimation. *IEEE Access*, 8:133330–133348, 2020. [1](#)
- [29] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 483–499. Springer, 2016. [2](#)
- [30] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 483–499. Springer, 2016. [5](#)
- [31] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. [5](#)
- [32] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Ordinal depth supervision for 3d human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7307–7316, 2018. [2](#)
- [33] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7025–7034, 2017. [2](#)
- [34] Xiaoye Qian, Youbao Tang, Ning Zhang, Mei Han, Jing Xiao, Ming-Chun Huang, and Ruei-Sung Lin. HSTFormer: Hierarchical spatial-temporal transformers for 3d human pose estimation. *arXiv preprint arXiv:2301.07322*, 2023. [2](#), [6](#), [7](#)
- [35] N Dinesh Reddy, Laurent Guigues, Leonid Pishchulin, Jayan Eledath, and Srinivasa G Narasimhan. TesseTrack: End-to-end learnable multi-person articulated 3d pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15190–15200, 2021. [2](#)
- [36] Edoardo Remelli, Shangchen Han, Sina Honari, Pascal Fua, and Robert Wang. Lightweight multi-view 3d pose estimation through camera-disentangled representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6040–6049, 2020. [2](#)
- [37] 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 *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 461–478. Springer, 2022. [2](#), [6](#), [7](#)
- [38] Weijie Su, Xizhou Zhu, Chenxin Tao, Lewei Lu, Bin Li, Gao Huang, Yu Qiao, Xiaogang Wang, Jie Zhou, and Jifeng Dai. Towards all-in-one pre-training via maximizing multi-modal mutual information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15888–15899, 2023. [2](#)
- [39] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. [2](#), [6](#)
- [40] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European conference on computer vision (ECCV)*, pages 529–545, 2018. [2](#)
- [41] Zhenhua Tang, Zhaofan Qiu, Yanbin Hao, Richang Hong, and Ting Yao. 3D human pose estimation with spatio-temporal criss-cross attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4790–4799, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [3](#)
- [43] Jingbo Wang, Sijie Yan, Yuanjun Xiong, and Dahua Lin. Motion guided 3d pose estimation from videos. In *European Conference on Computer Vision*, pages 764–780. Springer, 2020. [3](#)
- [44] Julian Wiederer, Arij Bouazizi, Ulrich Kressel, and Vasileios Belagiannis. Traffic control gesture recognition for autonomous vehicles. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10676–10683. IEEE, 2020. [1](#)
- [45] Tianhan Xu and Wataru Takano. Graph stacked hourglass networks for 3d human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16105–16114, 2021. [2](#)
- [46] Bruce X. B. 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, 2023. [2](#), [3](#), [6](#), [7](#)
- [47] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. CoCa: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [2](#)
- [48] Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. MetaFormer is actually what you need for vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10819–10829, 2022. [3](#)
- [49] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Junsong Yuan. MixSTE: Seq2seq mixed spatio-temporal encoder for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13232–13242, June 2022. [2](#), [6](#), [7](#)
- [50] Zhe Zhang, Chunyu Wang, Weichao Qiu, Wenhui Qin, and Wenjun Zeng. Adafuse: Adaptive multiview fusion for accu-

- rate human pose estimation in the wild. *International Journal of Computer Vision*, 129:703–718, 2021. [2](#)
- [51] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3425–3435, 2019. [3](#)
- [52] Qitao Zhao, Ce Zheng, Mengyuan Liu, Pichao Wang, and Chen Chen. PoseFormerV2: Exploring frequency domain for efficient and robust 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8877–8886, June 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [53] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11656–11665, 2021. [2](#)
- [54] Kun Zhou, Xiaoguang Han, Nianjuan Jiang, Kui Jia, and Jiangbo Lu. HEMlets pose: Learning part-centric heatmap triplets for accurate 3d human pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2344–2353, 2019. [2](#)
- [55] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. MotionBERT: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#)
- [56] Zhuofan Zong, Guanglu Song, and Yu Liu. DETRs with collaborative hybrid assignments training. *arXiv preprint arXiv:2211.12860*, 2022. [2](#)