

DiffPose: SpatioTemporal Diffusion Model for Video-Based Human Pose Estimation

Runyang Feng^{1,2}, Yixing Gao^{1,2*}, Tze Ho Elden Tse³, Xueqing Ma^{1,2}, Hyung Jin Chang³

¹ School of Artificial Intelligence, Jilin University,

² Engineering Research Center of Knowledge-Driven Human-Machine Intelligence, Ministry of Education, China, ³ School of Computer Science, University of Birmingham

{fengry22, maxq21}@mails.jlu.edu.cn, gaoyixing@jlu.edu.cn,
txt994@student.bham.ac.uk, h.j.chang@bham.ac.uk

Abstract

Denosing diffusion probabilistic models that were initially proposed for realistic image generation have recently shown success in various perception tasks (e.g., object detection and image segmentation) and are increasingly gaining attention in computer vision. However, extending such models to multi-frame human pose estimation is non-trivial due to the presence of the additional temporal dimension in videos. More importantly, learning representations that focus on keypoint regions is crucial for accurate localization of human joints. Nevertheless, the adaptation of the diffusion-based methods remains unclear on how to achieve such objective. In this paper, we present DiffPose, a novel diffusion architecture that formulates video-based human pose estimation as a conditional heatmap generation problem. First, to better leverage temporal information, we propose SpatioTemporal Representation Learner which aggregates visual evidences across frames and uses the resulting features in each denoising step as a condition. In addition, we present a mechanism called Lookup-based Multi-Scale Feature Interaction that determines the correlations between local joints and global contexts across multiple scales. This mechanism generates delicate representations that focus on keypoint regions. Altogether, by extending diffusion models, we show two unique characteristics from DiffPose on pose estimation task: (i) the ability to combine multiple sets of pose estimates to improve prediction accuracy, particularly for challenging joints, and (ii) the ability to adjust the number of iterative steps for feature refinement without retraining the model. DiffPose sets new state-of-the-art results on three benchmarks: PoseTrack2017, PoseTrack2018, and PoseTrack21.

1. Introduction

Human pose estimation has been extensively studied in computer vision, with the aim of detecting all instances of

*Corresponding Author

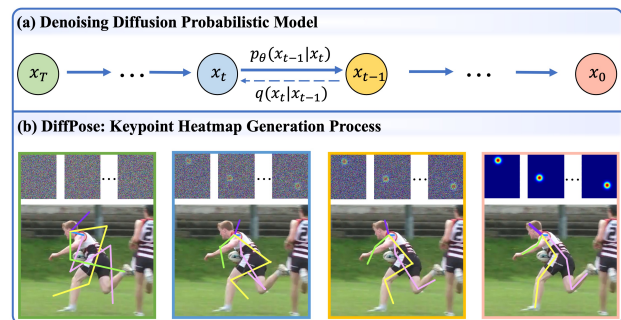


Figure 1. (a) Illustration of the original diffusion model where q and p_θ refer to the diffusion and denoising process, respectively. (b) In this work, we propose a novel framework named DiffPose which formulates video-based human pose estimation as a generative process of keypoint heatmaps.

people from images and localizing anatomical keypoints for each individual [20, 57, 63, 68]. It finds numerous applications ranging from human-computer interaction and augmented reality [21, 22] to behavior analysis and surveillance tracking [36, 40, 56, 64, 65, 66]. Conventional approaches [55, 69, 82] mainly employ the probabilistic graphical model or the pictorial structure model. Fueled by the explosion of deep learning, *Convolutional Neural Networks* [8, 39, 40, 63] and *Vision Transformers* [38, 75, 80] have witnessed significant progress in this task.

Until recently, denoising diffusion probabilistic models [30, 58], which are a type of generative models, have received much research attention for surpassing other methods such as GANs and achieving state-of-the-art generative results [6, 14]. The superior performance of the diffusion model has facilitated its expansion in diverse applications, such as super-resolution [54], inpainting [42], and image deblurring [52]. Following the demonstration of the effectiveness of diffusion models as representation learners for discriminative computer vision problems [6], several contemporary approaches have successfully employed the diffusion model for perception tasks, including object detec-

tion [10] and image segmentation [1, 6, 26].

Despite the considerable attention that diffusion models have gained following their achievements, their adaptation for video-based human pose estimation has significantly trailed that of other vision tasks, such as segmentation and object detection. We conjecture two primary reasons that underlie this disparity: (i) Effectively leveraging temporal information is crucial for video-based human pose estimation [39]. However, despite the success of various diffusion architectures in perception tasks, they are primarily designed for static images and incapable to capture temporal dependencies across frames. (ii) Real world images typically contain many task-irrelevant cues and accurately estimating human poses requires focusing on specific body joint regions [23]. However, it is still an open question on how to guide the diffusion model to filter out the unnecessary details and only attend to the keypoint regions.

In this paper, we present a novel architecture, termed SpatioTemporal **Diffusion Model for Pose Estimation** (DiffPose). By extending the framework of diffusion model, DiffPose presents a new approach to video-based human pose estimation. Specifically, it reformulates this problem as a conditional generative task of keypoint-wise heatmaps, as illustrated in Fig. 1. DiffPose consists of two primary stages: a forward diffusion stage that gradually introduces Gaussian noise to the ground truth heatmaps, and a reverse denoising stage that utilizes a Pose-Decoder to recover the original heatmap from the noisy input progressively.

Unlike the vanilla diffusion model [30], which simply uses U-Net [53] for denoising, we propose two novel designs that enhance the capabilities of the Pose-Decoder. These modifications enable the Pose-Decoder to better *utilize temporal information* and *focus on joint regions*. (i) We design a SpatioTemporal Representation Learner (STRL) which sequentially performs spatial information extraction within each frame and integrates cross-frame knowledge through cascaded Transformers. The resulting features, which contain rich temporal priors, are subsequently utilized as a fixed condition at each denoising step by the Pose-Decoder. (ii) In addition, we propose a Lookup-based MultiScale Feature Interaction mechanism (LMSFI), which guides the Pose-Decoder to learn intricate representations for pose prediction by inductively leveraging information from both noisy heatmaps and spatiotemporal features. To be specific, we first construct probabilistic joint fields based on the noisy heatmaps, and perform lookups over spatiotemporal features accordingly to activate keypoint region features. Then, we model fine-grained correlations between the retrieved local joint features and original global contexts over multiple scales to produce the final representations. By conducting feature interaction through LMSFI, we can explicitly reason about the relationships between joints and global contexts. As shown in Figure 4, our

proposed method is able to learn representations that converge around keypoint regions consistently.

An important feature of the diffusion-based framework is the separation of model training and evaluation. To provide more context, DiffPose is trained to reverse the forward diffusion process (*i.e.*, predict ground truth heatmaps from noises) and performs multi-step denoising to generate predictions based on randomly generated noisy heatmaps at inference. Benefiting from such framework, we demonstrate two distinct properties that appeal to human pose estimation task. (i) As DiffPose can generate multiple plausible pose estimates by sampling random noises, they can be combined to improve the prediction robustness, especially for challenging joints such as wrists and ankles. (ii) In contrast to existing methods that adopt a fixed iterative refinement structure [9, 71, 45], DiffPose can adaptively vary the number of denoising steps without retraining the model. From extensive experiments, we show that DiffPose consistently outperforms existing well-established approaches on three benchmark datasets. Furthermore, each of our proposed design choices is verified through ablation studies.

The key contributions of this work are summarized as follows: (1) To our best knowledge, we are the first to investigate video-based human pose estimation from the lens of generative modeling. In particular, we propose DiffPose, the first model that applies diffusion model to multi-frame human pose estimation. (2) We demonstrate two properties of DiffPose that are effective on pose estimation: the ability to enhance performance by aggregating multiple pose estimations and to perform flexible iterative refinement without model retraining. (3) We show that our DiffPose delivers state-of-the-art results on three benchmark datasets, PoseTrack2017, PoseTrack2018, and PoseTrack21.

2. Related Work

Human pose estimation. Early efforts on human pose estimation focus on static images, starting from building probability graphical structures [35, 28] to model the relations between body joints. With the advancement of deep learning [29, 67] and the availability of large-scale benchmark datasets [33, 2, 15], various deep architectures (*e.g.*, CNN and Transformer)-based methods are currently the dominating solutions [72, 73, 63, 38, 80, 77, 75]. There are two mainstream paradigms: i) regressing the position of keypoints from image directly, and ii) estimating probability heatmaps to represent keypoints locations. Representation using heatmap has gained more popularity due to the performance derived from faster optimization convergence.

Conversely, various studies [60, 47, 39, 40, 70] have attempted to estimate human poses in videos. [39] merges heatmaps of consecutive frames and computes their residuals to obtain joint-level features for pose estimation. [40] performs implicit motion compensation using deformable

convolutions for better feature aggregation and heatmap prediction. These approaches usually predict a deterministic pose solution for each frame and lack effective recalibrations, which might suffer from localized detection failure especially for challenging joints. In contrast, Our approach benefiting from probabilistic diffusion models is able to combine multiple pose solutions naturally to provide more robust estimations.

Diffusion model. Diffusion models [30, 58] are a type of deep generative models that utilize the final state of a Markov chain originating from a standard Gaussian distribution to approximate the distribution of natural images. Neural network is typically trained to reverse this diffusion process for each Markov step. Within this framework, diffusion models have recently demonstrated remarkable results in a wide spectrum of generative tasks from visual images [54, 4, 26, 46, 18, 44] to nature language [3, 25, 37, 31, 48, 76]. Diffusion models have also proven to be useful in various discriminative computer vision problems [10, 1, 6, 26]. The pioneer work [1] presents a diffusion model conditioned on an input image for image segmentation. [10] proposes DiffusionDet which formulates object detection as a generative denoising process from noisy boxes to object boxes. [26] further extends [10] to perform instance segmentation. To the best of our knowledge, there have been no previous successful attempts to adapt diffusion models for multi-frame human pose estimation. This paper introduces DiffPose, which explores the potential of diffusion models in video-based human pose estimation and is the first diffusion model to achieve state-of-the-art performance for this task.

3. Our Approach

3.1. Preliminaries

Problem Formulation. Following the top-down pose estimation approach, we first obtain all human bounding boxes per frame I_t using an off-the-shelf object detector. Each bounding box is then enlarged by 25% to crop the same individual in consecutive frames $\mathcal{I}_t = \langle I_{t-\delta}, \dots, I_t, \dots, I_{t+\delta} \rangle$ with δ being a predefined temporal interval. In this way, we obtain the cropped video segment $\mathcal{I}_t^i = \langle I_{t-\delta}^i, \dots, I_t^i, \dots, I_{t+\delta}^i \rangle$ for person i . Given an image sequence \mathcal{I}_t^i centered on the key frame I_t^i , our goal is to estimate the keypoint heatmaps for I_t^i .

Diffusion Model. Inspired by non-equilibrium thermodynamics [61, 62], diffusion models are under the category of latent variable models which aim to reconstruct a task-specific distribution that starts from random noise. These models typically consist of two basic processes: 1) a forward process that gradually adds Gaussian noise to sample data, and 2) a reverse process that learns to invert the forward diffusion. To be specific, the forward diffusion pro-

cess is defined as:

$$\begin{aligned} q(\mathbf{x}_t | \mathbf{x}_0) &:= \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}), \\ \mathbf{x}_t &= \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \epsilon \sim \mathcal{N}(0, 1), \end{aligned} \quad (1)$$

where $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s = \prod_{s=1}^t (1 - \beta_s)$ and β_s denotes the noise variance schedule [30]. The operation in Eq. 1 adds noise to the original data sample \mathbf{x}_0 and transforms it into a latent noisy sample \mathbf{x}_t at an arbitrary sampling step $t \in \{0, 1, \dots, T\}$. During training, a neural network $f_\theta(\mathbf{x}_t, t)$ is trained to perform the denoising task either by predicting \mathbf{x}_0 or ϵ (we choose the former as done in [10, 11]), with the constraint of L_2 loss. This process is expressed as:

$$\mathcal{L}_{\mathbf{x}_0} = \|f_\theta(\mathbf{x}_t, t) - \mathbf{x}_0\|^2. \quad (2)$$

In inference, the learned denoising (reverse) function f_θ is applied to a random noise sample \mathbf{x}_T along with a preset updating rule [30, 59], to reconstruct the data sample \mathbf{x}_0 in an iterative way $\mathbf{x}_T \rightarrow \mathbf{x}_{T-\Delta} \rightarrow \dots \rightarrow \mathbf{x}_0$.

In this paper, we propose a novel framework that enables the diffusion model to better process dynamic contexts for video-based human pose estimation. Specifically, we present DiffPose which modulates the vanilla diffusion model to *incorporate temporal information* and *attend to keypoint region cues*, resulting in a paradigm more aligned with multi-frame human pose estimation. Our proposed framework is illustrated in Fig. 2. In our problem setting, the original data sample is the ground truth heatmap $\mathbf{x}_0 = \mathbf{H}_t^i$. This heatmap is generated using a 2D Gaussian centered at the annotated joint location. We train a Pose-Decoder $f_\theta(\mathbf{x}_t, \mathbb{F}_t^i, t)$ to recover \mathbf{x}_0 from the noisy heatmap \mathbf{x}_t by conditioning on the spatiotemporal feature of the input sequence \mathbb{F}_t^i which is derived by the SpatioTemporal Representation Learner (STRL).

In the following, we first detail the architecture of STRL (Sec. 3.2) and Pose-Decoder (Sec. 3.3). Then, we explain the training and inference algorithms (Sec. 3.4) as well as providing discussions on the favorable properties of DiffPose for pose estimation in Sec. 3.5.

3.2. Spatiotemporal Representation Learner

Inspired by the success of Vision Transformers [17, 75, 41], we employ cascaded Transformers to capture the spatial-temporal dependencies among video frames. Given the sequence data $\mathcal{I}_t^i = \langle I_{t-\delta}^i, \dots, I_t^i, \dots, I_{t+\delta}^i \rangle$ as input, we first employ a plain Vision Transformer [17, 75] pretrained on ImageNet [13] as the backbone network to extract spatial features $\langle F_{t-\delta}^i, \dots, F_t^i, \dots, F_{t+\delta}^i \rangle$ for each frame. Subsequently, each frame feature is spatially rearranged and fed into a patch embedding layer, which embeds the feature into tokens $\langle \bar{F}_{t-\delta}^i, \dots, \bar{F}_t^i, \dots, \bar{F}_{t+\delta}^i \rangle$. Then, we concatenate all embedded tokens, retain their spatial information through a learnable position embedding E_{POS} , and feed them into

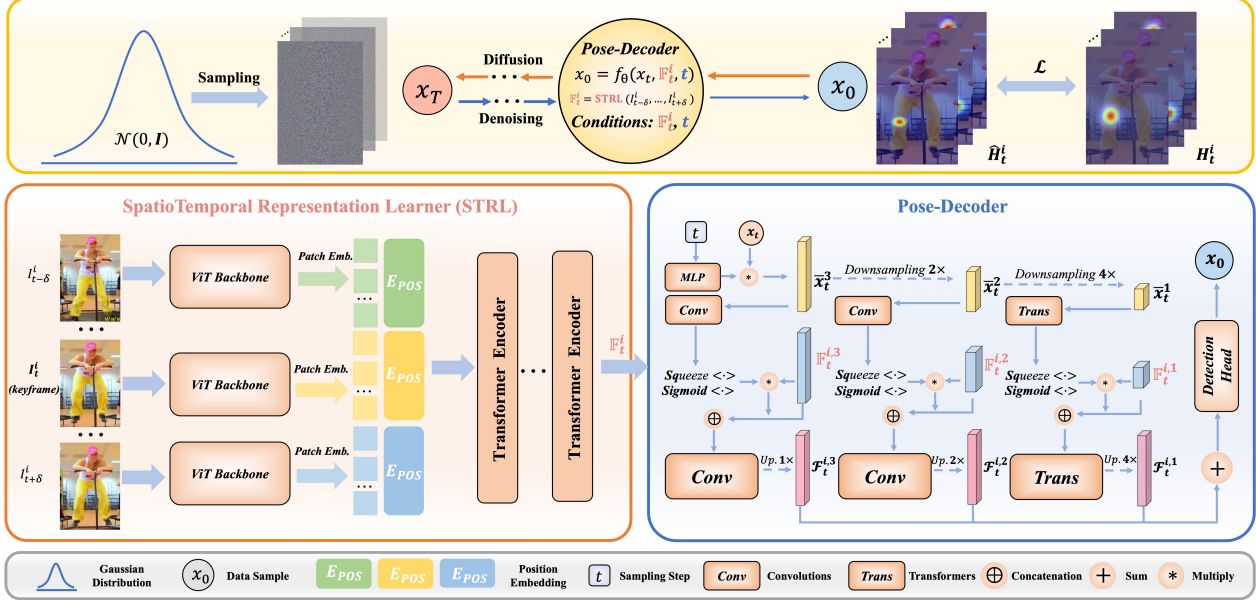


Figure 2. Overall pipeline of the proposed DiffPose framework. The goal is to detect the human pose in the keyframe I_t^i . Given an input sequence, our SpatioTemporal Representation Learner (STRL) extracts the spatiotemporal feature \mathbb{F}_t^i . The feature \mathbb{F}_t^i , the noisy heatmap \mathbf{x}_t , and the sampling step t are then passed to the Pose-Decoder, which performs lookup-based multiscale feature interaction to obtain representations $\mathcal{F}_t^i = \{\mathcal{F}_t^{i,1}, \mathcal{F}_t^{i,2}, \mathcal{F}_t^{i,3}\}$. Finally, these features are aggregated to attain the final pose estimation $\hat{\mathbf{H}}_t^i$ (i.e., \mathbf{x}_0).

cascaded Transformer encoders where each encoder consists of a Multi-Head Self-Attention (MHSA) layer and a feed-forward neural network (FFN). Finally, the encoded deep features of all frames are aggregated via a Multilayer Perceptron (MLP) to produce the spatiotemporal feature \mathbb{F}_t^i . The above procedures can be formulated as:

$$\begin{aligned}
 \tilde{F}_t^0 &= \text{Concat}(\bar{F}_{t-\delta}^i + E_{POS}^{t-\delta}, \dots, \bar{F}_{t+\delta}^i + E_{POS}^{t+\delta}), \\
 \tilde{F}_t^l &= \tilde{F}_t^{l-1} + \text{MHSA}(\text{LN}(\tilde{F}_t^{l-1})), \\
 \tilde{F}_t^l &= \tilde{F}_t^l + \text{FFN}(\text{LN}(\tilde{F}_t^l)), \\
 &\vdots \\
 \mathbb{F}_t^i &= \text{MLP}(\text{LN}(\tilde{F}_t^L)),
 \end{aligned} \tag{3}$$

where the superscript $l \in [1, 2, \dots, L]$ denotes the output of l -th Transformer layer and \tilde{F}_t^0 represents the initial feature. The function $\text{LN}(\cdot)$ indicates the LayerNorm layer. Note that the spatial (i.e., the number of tokens) and channel dimensions within each transformer layer remain constant.

3.3. Pose-Decoder

After obtaining the spatiotemporal feature \mathbb{F}_t^i , the Pose-Decoder denoises the heatmap \mathbf{x}_t by taking \mathbb{F}_t^i together with the sampling step t as conditions, and output the predicted heatmap $\hat{\mathbf{x}}_0 = \hat{\mathbf{H}}_t^i$. Specifically, we first project the step index t into an embedding and utilize the embedding to rescale the initial noisy heatmap \mathbf{x}_t , attaining the step-adaptive version $\bar{\mathbf{x}}_t$. Then, we model the global correlations between \mathbb{F}_t^i and $\bar{\mathbf{x}}_t$ across multiple scales via Transformer or

convolutional structures, and obtain multi-scale representations $\mathcal{F}_t^i = \{\mathcal{F}_t^{i,1}, \mathcal{F}_t^{i,2}, \mathcal{F}_t^{i,3}\}$. Finally, these features are integrated and passed to a detection head to predict the pose heatmap $\hat{\mathbf{H}}_t^i$.

In order to encourage the representations \mathcal{F}_t^i to focus on keypoint regions, we propose a Lookup-based Multi-Scale Feature Interaction mechanism (LMSFI) to inductively model correlations between \mathbb{F}_t^i and $\bar{\mathbf{x}}_t$. It consists of two procedures: pairwise size-matched feature generation and lookup-based feature interaction.

Pairwise size-matched feature generation. Given $\mathbb{F}_t^i \in \mathbb{R}^{C \times H \times W}$ and $\bar{\mathbf{x}}_t \in \mathbb{R}^{c \times 4H \times 4W}$ with different spatial dimensions, we perform upsampling and downsampling separately to construct size-matched feature pairs $\langle \mathbb{F}_t^i, \bar{\mathbf{x}}_t \rangle$. Specifically, we adopt several deconvolution layers to perform $1\times$, $2\times$, and $4\times$ upsampling of resolution over \mathbb{F}_t^i , and obtain corresponding features $\{\mathbb{F}_t^{i,1}, \mathbb{F}_t^{i,2}, \mathbb{F}_t^{i,3}\}$. Similarly, stride convolutions are used to downsample on $\bar{\mathbf{x}}_t$ to produce $\{\bar{\mathbf{x}}_t^1, \bar{\mathbf{x}}_t^2, \bar{\mathbf{x}}_t^3\}$ ($\bar{\mathbf{x}}_t = \bar{\mathbf{x}}_t^3$). With the above process, we attain multi-scale size-matched feature pairs $\langle \mathbb{F}_t^{i,J}, \bar{\mathbf{x}}_t^J \rangle$. The superscript $J = \{1, 2, 3\}$ refers to resolutions of different levels from low to high.

Lookup-based feature interaction. Upon constructing multi-scale feature pairs $\langle \mathbb{F}_t^{i,J}, \bar{\mathbf{x}}_t^J \rangle$, we model interactions between the spatiotemporal feature $\mathbb{F}_t^{i,j}$ and the noisy heatmap $\bar{\mathbf{x}}_t^j$ at each resolution j individually to obtain corresponding feature representation $\mathcal{F}_t^{i,j}$. A naive approach would be to *directly* concatenate and aggregate $\mathbb{F}_t^{i,j}$ and $\bar{\mathbf{x}}_t^j$. In our experiments, we show that the learned keypoint features of this scheme are scattered across significant areas

(see Fig. 4), resulting in performance reduction as shown in Table 4. In practice, heatmaps reveal the likelihood of the locations containing joints, whereas the noisy heatmap $\bar{\mathbf{x}}_t^j$ is corrupted and can provide a negligible amount of valid real-valued information [12]. As a result, directly modeling correlations of $\bar{\mathbf{x}}_t^j$ and $\mathbb{F}_t^{i,j}$ is extremely challenging. Therefore, we adopt an inductive modeling strategy which first performs lookups over the spatiotemporal feature $\mathbb{F}_t^{i,j}$ according to the heatmap $\bar{\mathbf{x}}_t^j$ to retrieve local joint feature $\bar{\mathbb{F}}_t^{i,j}$, and then models correlations between the local feature $\bar{\mathbb{F}}_t^{i,j}$ and the vanilla global context $\mathbb{F}_t^{i,j}$.

More specifically, considering that the computational complexity of self-attention increases quadratically with input resolution, we adopt a composite structure that uses Transformers and convolutions to capture feature interactions at low and high resolutions, respectively. **(i)** For the low resolution feature pair $\langle \mathbb{F}_t^{i,1}, \bar{\mathbf{x}}_t^1 \rangle$, the noisy heatmap $\bar{\mathbf{x}}_t^1$ is first embedded to feature tokens, and a Transformer encoder is leveraged to perform self-refinement to yield $\bar{\bar{\mathbf{x}}}_t^1$. Then, we take the maximum activations along the depth dimension over $\bar{\bar{\mathbf{x}}}_t^1$ to squeeze the global channel information into a single-channel descriptor, followed by a sigmoid function to obtain an attention mask A^1 that indicates possible keypoint fields. Then, the mask A^1 is used to retrieve corresponding spatiotemporal features $\bar{\mathbb{F}}_t^{i,1}$. Finally, $\bar{\mathbb{F}}_t^{i,1}$ and $\mathbb{F}_t^{i,1}$ are concatenated and processed by cascaded Transformers, followed by upsampling of resolution to output feature $\mathcal{F}_t^{i,1}$. The above process can be described as:

$$\begin{aligned} \bar{\bar{\mathbf{x}}}_t^1 &= \text{SeRef}(\bar{\mathbf{x}}_t^1), & A^1 &= \text{Sigmoid}(\text{Sq}(\bar{\bar{\mathbf{x}}}_t^1)), \\ \bar{\mathbb{F}}_t^{i,1} &= A^1 \odot \mathbb{F}_t^{i,1}, & \mathcal{F}_t^{i,1} &= \text{Up}(\text{Trans}(\bar{\mathbb{F}}_t^{i,1} \oplus \mathbb{F}_t^{i,1})), \end{aligned} \quad (4)$$

where $\text{SeRef}(\cdot)$, $\text{Sq}(\cdot)$, \odot , \oplus , and $\text{Up}(\cdot)$ denote the operations of self-refinement, squeezing, spatial-wise multiplication, concatenation, and upsampling, respectively. **(ii)** For high-resolution feature pairs $\langle \mathbb{F}_t^{i,j}, \bar{\mathbf{x}}_t^j \rangle$ with $j = 2, 3$, an analogical procedure is executed using convolutions:

$$\begin{aligned} \bar{\bar{\mathbf{x}}}_t^j &= \text{Conv}(\bar{\mathbf{x}}_t^j), & A^j &= \text{Sigmoid}(\text{Sq}(\bar{\bar{\mathbf{x}}}_t^j)), \\ \bar{\mathbb{F}}_t^{i,j} &= A^j \odot \mathbb{F}_t^{i,j}, & \mathcal{F}_t^{i,j} &= \text{Up}(\text{Conv}(\bar{\mathbb{F}}_t^{i,j} \oplus \mathbb{F}_t^{i,j})). \end{aligned} \quad (5)$$

Heatmap generation. Ultimately, we integrate feature representations across all scales $\{\mathcal{F}_t^{i,1}, \mathcal{F}_t^{i,2}, \mathcal{F}_t^{i,3}\}$ via element-wise addition, and employ a detection head (*i.e.*, a 3×3 convolution) to yield the predicted heatmap $\hat{\mathbf{H}}_t^i$. By inductively modeling multi-scale feature interactions, our Pose-Decoder is able to reason about the fine-grained relations of keypoints and global contexts, thereby producing more tailored representations that attend to joint areas.

3.4. Overall Training and Inference Algorithms

Training. We perform diffusion process that corrupts ground truth heatmaps to noisy heatmaps, and train the

Pose-Decoder for heatmap denoising to reverse this process. Specifically, we sample Gaussian noises according to α_t in Eq. 1 and add them to ground truth heatmaps to obtain the noisy samples. The parameter α_t at each sampling step t is predefined by a monotonically decreasing cosine scheme, as adopted in [30]. We employ the standard pose estimation loss (mean square error) to supervise the model training:

$$\mathcal{L} = \left\| \mathbf{H}_t^i - \hat{\mathbf{H}}_t^i \right\|_2^2, \quad (6)$$

where \mathbf{H}_t^i and $\hat{\mathbf{H}}_t^i$ denote the ground truth and predicted pose heatmaps, respectively.

Inference. The proposed DiffPose conducts denoising on noisy heatmaps sampled from Gaussian distribution, progressively refining its predictions over multiple sampling steps. For each sampling step, the Pose-Decoder takes random noisy heatmaps or the predicted heatmaps of the last sampling step as input and outputs the estimated heatmaps of the current step. Then, we adopt DDIM [59] to update the heatmaps for the next step.

3.5. Discussion

Building on diffusion-based architecture, our proposed framework DiffPose is able to decouple training and testing stages which enables a more adaptable inference process. By extending this concept, we investigate further and demonstrate the unique benefits of DiffPose for pose estimation, specifically in the areas of *flexible pose ensemble* and *flexible iterative refinement*.

Flexible pose ensemble. In common practices, one usually performs inference starting with a single initial sample. However, the diffusion model is intrinsically probabilistic [1] and can generate diverse outputs for different noise inputs. Correspondingly, taking different noisy heatmaps as input, the DiffPose can yield different plausible pose predictions that possess respective keypoint recognition preferences. Ensembling these complementary pose solutions can enhance the robustness and stability of model predictions especially for challenging joints. To exploit this phenomenon, we initialize N groups of noisy heatmaps for inference and subsequently average their predictions. Experimental results show that the complementary pose ensemble brings significant performance improvement (see Table 6).

Flexible iterative refinement. After training the model, the DiffPose performs multi-step refinement (sampling) progressively to yield the final pose prediction. In practice, the number of sampling steps can be adjusted flexibly without retraining the model, which is preferable to the prior approaches that adopt a fixed structure of iterative refinement [9, 71, 43, 45]. By increasing iterative sampling steps, the resulting representations would be more delicate, which fosters accurate pose estimation (Fig. 4).

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
PoseTracker [24]	67.5	70.2	62.0	51.7	60.7	58.7	49.8	60.6
PoseFlow [74]	66.7	73.3	68.3	61.1	67.5	67.0	61.3	66.5
JointFlow [16]	-	-	-	-	-	-	-	69.3
FastPose [81]	80.0	80.3	69.5	59.1	71.4	67.5	59.4	70.3
TML++ [32]	-	-	-	-	-	-	-	71.5
Simple (R-50) [73]	79.1	80.5	75.5	66.0	70.8	70.0	61.7	72.4
Simple (R-152) [73]	81.7	83.4	80.0	72.4	75.3	74.8	67.1	76.7
STEmbedding [34]	83.8	81.6	77.1	70.0	77.4	74.5	70.8	77.0
HRNet [63]	82.1	83.6	80.4	73.3	75.5	75.3	68.5	77.3
MDPN [27]	85.2	88.5	83.9	77.5	79.0	77.0	71.4	80.7
CorrTrack [50]	86.1	87.0	83.4	76.4	77.3	79.2	73.3	80.8
Dynamic-GNN [78]	88.4	88.4	82.0	74.5	79.1	78.3	73.1	81.1
PoseWarper [8]	81.4	88.3	83.9	78.0	82.4	80.5	73.6	81.2
DCPose [39]	88.0	88.7	84.1	78.4	83.0	81.4	74.2	82.8
DetTrack [70]	89.4	89.7	85.5	79.5	82.4	80.8	76.4	83.8
FAMI-Pose [40]	89.6	90.1	86.3	80.0	84.6	83.4	77.0	84.8
DiffPose (Ours)	89.0	91.2	87.4	83.5	85.5	87.2	80.2	86.4

Table 1. Quantitative results on the **PoseTrack2017** validation set.

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
STAF [49]	-	-	-	64.7	-	-	62.0	70.4
AlphaPose [19]	63.9	78.7	77.4	71.0	73.7	73.0	69.7	71.9
TML++ [32]	-	-	-	-	-	-	-	74.6
MDPN [27]	75.4	81.2	79.0	74.1	72.4	73.0	69.9	75.0
PGPT [5]	-	-	-	72.3	-	-	72.2	76.8
Dynamic-GNN [78]	80.6	84.5	80.6	74.4	75.0	76.7	71.8	77.9
PoseWarper [8]	79.9	86.3	82.4	77.5	79.8	78.8	73.2	79.7
PT-CPN++ [79]	82.4	88.8	86.2	79.4	72.0	80.6	76.2	80.9
DCPose [39]	84.0	86.6	82.7	78.0	80.4	79.3	73.8	80.9
DetTrack [70]	84.9	87.4	84.8	79.2	77.6	79.7	75.3	81.5
FAMI-Pose [40]	85.5	87.7	84.2	79.2	81.4	81.1	74.9	82.2
DiffPose (Ours)	85.0	87.7	84.3	81.5	81.4	82.9	77.6	83.0

Table 2. Quantitative results on the **PoseTrack2018** validation set.

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
Tracker++ w. poses [7, 15]	-	-	-	-	-	-	-	71.4
CorrTrack [50, 15]	-	-	-	-	-	-	-	72.3
CorrTrack w. ReID [50, 15]	-	-	-	-	-	-	-	72.7
Tracker++ w. corr. [7, 15]	-	-	-	-	-	-	-	73.6
DCPose [39]	83.2	84.7	82.3	78.1	80.3	79.2	73.5	80.5
FAMI-Pose [40]	83.3	85.4	82.9	78.6	81.3	80.5	75.3	81.2
DiffPose (Ours)	84.7	85.6	83.6	80.8	81.4	83.5	80.0	82.9

Table 3. Quantitative results on the **PoseTrack21** dataset.

4. Experiments

4.1. Experimental Settings

Datasets. We benchmark the proposed DiffPose on three widely used benchmark datasets for video-based human pose estimation, PoseTrack2017 [33], PoseTrack2018 [2], and PoseTrack21 [15]. These datasets contain video sequences of tricky scenarios where clustered people perform rapid movement. Specifically, **PoseTrack2017** includes 250 video clips for training and 50 videos for validation, with a total of 80,144 pose annotations. **PoseTrack2018** considerably increases the number of clips, containing 593 videos for training, 170 videos for validation, and a total of 153,615 pose annotations. Both datasets identify 15 keypoints, with an additional label for joint visibility. The training videos are densely annotated in the center 30 frames, and validation videos are additionally labeled every four frames. **PoseTrack21** further enriches PoseTrack2018 especially for annotations of small persons and persons in crowds, including 177,164 human pose annotations.

Evaluation metric. The performance of the proposed method is evaluated with the widely-adopted [70, 63, 8, 39] human pose estimation metric namely average precision (**AP**). We compute the AP for each joint and then average overall joints to obtain the final performance (**mAP**).

Implementation details. Our DiffPose is implemented with PyTorch. The input image size is fixed to 256×192 . We incorporate data augmentation including random rotation $[-45^\circ, 45^\circ]$, random scale $[0.65, 1.35]$, truncation (half body), and flipping during training. The time interval δ is set to 2. We define the total sampling steps $T = 1000$. We adopt the AdamW [51] optimizer with a base learning rate of $5e-4$ (decays to $5e-5$ and $5e-6$ at the 20^{th} and 40^{th} epochs, respectively). We train the model using 4 TITAN RTX GPUs. All training process is terminated within 60 epochs. During inference, we initialize $N = 10$ groups of noises, and set the iterative denoising steps to 4.

4.2. Comparison with State-of-the-art Approaches

Results on the PoseTrack2017 dataset. We first benchmark our model on the PoseTrack2017 dataset. A total of 17 methods are compared, and their performances on the validation set are summarized in Table 1. We can observe that our DiffPose delivers state-of-the-art pose estimation performance compared to existing well-established approaches, by adopting a generative paradigm for the first time. DiffPose attains the final performance of 86.4 mAP, and provides a 1.6 mAP gain over the previous best-performed method FAMI-Pose [40]. The performance boost for challenging joints (*i.e.*, wrist, ankle) is also encouraging: we achieve an mAP of 83.5 ($\uparrow 3.5$) for wrists and an mAP of 80.2 ($\uparrow 3.2$) for ankles. Such consistent performance improvements suggest the great potential of diffusion models in pose estimation. Another observation is that pose estimation methods that integrate temporal information (such as DetTrack and FAMI-Pose) indeed surpass approaches using only the single keyframe. This corroborates the importance of our design that injects spatiotemporal features into the DiffPose model. Furthermore, we show example visualization results in Fig. 3, which are indicative of the robustness of our method in tricky scenes.

Results on the PoseTrack2018 dataset. We further evaluate the proposed DiffPose on the PoseTrack2018 dataset, and report the detailed results on the validation set in Table 2. As shown in this table, our DiffPose once again outperforms all approaches and delivers the best results. We obtain the final performance of 83.0 mAP, with an mAP of 84.3, 81.5, and 77.6 for the elbow, wrist, and ankle joints.

Results on the PoseTrack21 dataset. Performance comparisons of our model and previous SOTA methods on the PoseTrack21 dataset are provided in Table 3. We observe that existing method FAMI-Pose [40] has already achieved an impressive performance of 81.2 mAP. In con-



Figure 3. Visual results of our DiffPose on benchmarks. Challenging scenarios including fast motion and mutual occlusion are involved.

Method	Spatiotemporal. (STRL)	Lookup. (LMSFI)	Mean
Direct-single			52.5
(a)		✓	82.4
(b)	✓	✓	86.4

Table 4. Ablation of different components in DiffPose.

Method	Multi-scale feature	Aggregation	Mean
Low-resolution			84.4
(a)	✓	✓	86.4
(b)	✓	✓*	86.2

Table 5. Ablation of various designs in LMSFI. * denotes fusing multi-scale features with concatenation and aggregation.

trast, DiffPose is able to achieve 82.9 ($\uparrow 1.7$) mAP. We also obtain an mAP of 80.8 for wrists and 80.0 for ankles.

4.3. Ablation Study

We perform ablation experiments to investigate the contribution of each component in our DiffPose, including SpatioTemporal Representation Learner (STRL) as well as the Lookup-based MultiScale Feature Interaction mechanism (LMSFI). We also examine the efficacy of various design choices in LMSFI. Finally, we study the influence of modifying the number of initial noise heatmaps and iterative steps (during inference) on the final performance.

Study on components of DiffPose. We empirically evaluate the effectiveness of each proposed component, and report the results in Table 4. We first construct a simple baseline namely Direct-single, which takes the single keyframe as condition and directly concatenates and aggregates noisy heatmap with image features to model their interactions. This straightforward scheme produces a severely degraded pose estimation performance of 52.5 mAP. This is in line with our intuitions, *i.e.*, the noisy heatmap is corrupted and usually contains distracting information, which incurs inherent difficulties in directly learning the correlations between noisy heatmaps and image features. **(a)** For this setting, we introduce the proposed LMSFI into the baseline Direct-single. Remarkably, the LMSFI inductively models feature interactions and improves performance from an mAP of 52.5 to 82.4. This significant performance boost

($\uparrow 29.9$ mAP) corroborates the importance of our LMSFI in guiding the model to focus on specific joint regions. **(b)** The final setting further incorporates spatiotemporal cues as the sampling condition and corresponds to our full DiffPose. The performance improvement of 4.0 mAP demonstrates the effectiveness of our DiffPose in introducing temporal information to facilitate video-based pose estimation.

Study on Lookup-based MultiScale Feature Interaction. We further explore the influence of various designs within LMSFI, and tabulate the results in Table 5. We model feature interactions in a single low resolution to form the baseline Low-resolution. **(a)** We introduce multi-scale fusion to the baseline method and produce the final performance of 86.4 ($\uparrow 2.0$) mAP. This significant performance improvement upon incorporation of the multi-scale fusion highlights its effectiveness in learning informative representations for better performance. **(b)** We also examine the impact of fusing multi-scale features by concatenation and aggregation. The result in mAP (86.2) changes marginally.

Study on initial noises. As discussed in Sec. 3.5, we propose a pose ensemble strategy that randomly initializes N groups of Gaussian noises during inference and averages their predictions. Table 6 shows the effects of adopting different N , where N is set to 1, 5, and 10. The quantitative results in mAP reflect a gradual performance improvement with increasing initial noises, from 80.4 \rightarrow 84.0 \rightarrow 86.4. This phenomenon can be attributed to the probabilistic nature of the diffusion model, whereby DiffPose is able to forecast diverse plausible poses from different noises. Ensembling such pose solutions enhances the robustness of model predictions and significantly boosts the pose estimation performance. Another observation is that the improvement in mAP with increasing N mainly stems from challenging joints such as wrists ($\uparrow 8.9$ mAP) and ankles ($\uparrow 6.1$ mAP), and this fact still remains valid compared to FAMI-Pose [40]. This suggests that the proposed pose ensemble strategy derived from the diffusion-based architecture can potentially facilitate the pose detection in intractable scenes.

Study on denoising steps. As discussed in Sec. 3.5, DiffPose can adopt an arbitrary number of iterative sam-

Method	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
FAMI-Pose [40]	89.6	90.1	86.3	80.0	84.6	83.4	77.0	84.8
$N = 1$	85.4	87.0	80.1	74.6	78.1	82.6	74.1	80.4
$N = 5$	87.4	88.8	83.7	79.6	83.7	84.8	78.1	84.0
$N = 10$, DiffPose	89.0	91.2	87.4	83.5	85.5	87.2	80.2	86.4

Table 6. Ablation of modifying the number of initial noises N .

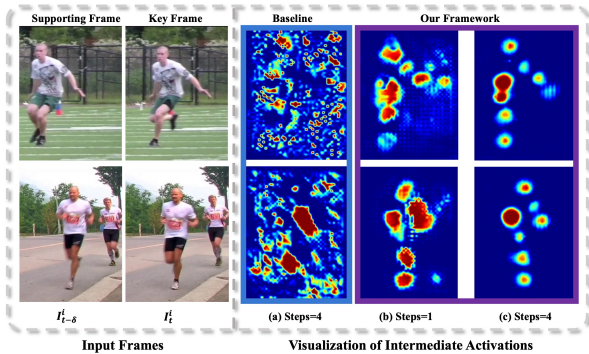


Figure 4. Visualizations of intermediate activations of the Direct-multi baseline (a) and our DiffPose at different denoising steps (b) $Steps = 1$ and (c) $Steps = 4$.

pling steps. To investigate how the number of iterative steps affects the final performance, we experiment with $Steps \in \{1, 2, 4\}$ and report the results in Table 7. It is clear that more iteration steps result in better performances. This is in accordance with our expectations, *i.e.*, the captured features are progressively refined to focus on keypoint regions upon multiple iterations, leading to better results.

4.4. Qualitative Analyses on DiffPose

Representation visualization. In addition to the quantitative results, we also provide qualitative analyses to better understand the mechanism behind DiffPose. Fig. 4 displays the intermediate activations of the Direct-multi baseline (incorporating STRL into the Direct-single) as well as our DiffPose. We observe that the DiffPose produces compact representations (b) and (c) that attend to local keypoint regions, while the features derived from the baseline (a) spread across salient areas. This provides empirical evidence that our LMSFI is effective in learning tailored representations for pose estimation. On the other hand, the features upon multi-step refinement are more attentive that encompass less task-irrelevant information.

Visual comparisons. We further examine the ability of our model in dealing with challenging scenarios such as mutual occlusion and fast motion. We depict in Fig. 5 the side-by-side comparisons of a) our DiffPose against state-of-the-art methods b) FAMI-Pose [40] and c) HRNet [63]. It is observed that our DiffPose consistently produces more robust pose predictions for various challenging scenes. HRNet [63] is designed for static images and does not incorporate temporal cues, leading to suboptimal results in degraded frames. On the other hand, FAMI-Pose [40] adopts a deterministic pose estimation paradigm, yielding a single pose solution for each person. Through the principled de-

Sampling Steps	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Mean
$Steps = 1$	88.9	90.4	85.9	82.5	85.1	86.6	80.2	85.9
$Steps = 2$	89.0	91.0	87.0	83.4	85.4	87.1	80.2	86.3
$Steps = 4$, DiffPose	89.0	91.2	87.4	83.5	85.5	87.2	80.2	86.4

Table 7. Ablation of modifying the number of denoising steps.



Figure 5. Visual comparisons of the pose estimation results of our DiffPose (a), FAMI-Pose (b), and HRNet-W48 (c) on the challenging cases from PoseTrack dataset. Inaccurate detections are highlighted by the red circles.

sign of the model architecture (*i.e.*, STRL, LMSFI) as well as the flexible training and testing pipeline of the diffusion model, our DiffPose is more adept at handling tricky cases.

5. Conclusion and Future Works

In this paper, we explore the video-based human pose estimation task from the perspective of generative modeling. We present a novel framework termed DiffPose which treats multi-frame human pose estimation as a conditional generative process of keypoint heatmaps. We design a SpatioTemporal Representation Learner (STRL) to integrate temporal clues into the diffusion model, as well as a Lookup-based MultiScale Feature Interaction mechanism (LMSFI) for inducing the model to attend to keypoint regions. Furthermore, we show two attractive properties of DiffPose for pose estimation including flexible pose ensemble and iterative refinement, which enable enhanced performance without retraining the model. Empirical evaluations on three standard benchmark datasets, PoseTrack2017, PoseTrack2018, and PoseTrack21 demonstrate that DiffPose achieves state-of-the-art performance. Future works include applying DiffPose to other vision tasks such as 3D human pose estimation and pose tracking.

6. Acknowledgements

This work is supported in part by the National Natural Science Foundation of China under Grant No. 62203184 and the International Cooperation Project under Grant No. 20220402009GH. This work is also supported in part by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2023-2020-0-01789), supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation).

References

- [1] Tomer Amit, Eliya Nachmani, Tal Shaharbany, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. [2](#), [3](#), [5](#)
- [2] Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. Posetrack: A benchmark for human pose estimation and tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#), [6](#)
- [3] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021. [3](#)
- [4] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022. [3](#)
- [5] Qian Bao, Wu Liu, Yuhao Cheng, Boyan Zhou, and Tao Mei. Pose-guided tracking-by-detection: Robust multi-person pose tracking. *IEEE Transactions on Multimedia*, 23:161–175, 2020. [6](#)
- [6] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khulkov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. [1](#), [2](#), [3](#)
- [7] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 941–951, 2019. [6](#)
- [8] Gedas Bertasius, Christoph Feichtenhofer, Du Tran, Jianbo Shi, and Lorenzo Torresani. Learning temporal pose estimation from sparsely-labeled videos. In *Advances in Neural Information Processing Systems*, pages 3027–3038, 2019. [1](#), [6](#)
- [9] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. [2](#), [5](#)
- [10] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. [2](#), [3](#)
- [11] Ting Chen, Lala Li, Saurabh Saxena, Geoffrey Hinton, and David J Fleet. A generalist framework for panoptic segmentation of images and videos. *arXiv preprint arXiv:2210.06366*, 2022. [3](#)
- [12] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022. [5](#)
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [3](#)
- [14] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [1](#)
- [15] Andreas Doering, Di Chen, Shanshan Zhang, Bernt Schiele, and Juergen Gall. Posetrack21: A dataset for person search, multi-object tracking and multi-person pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20963–20972, 2022. [2](#), [6](#)
- [16] Andreas Doering, Umar Iqbal, and Juergen Gall. Joint flow: Temporal flow fields for multi person tracking. *arXiv preprint arXiv:1805.04596*, 2018. [6](#)
- [17] 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. [3](#)
- [18] Wan-Cyuan Fan, Yen-Chun Chen, DongDong Chen, Yu Cheng, Lu Yuan, and Yu-Chiang Frank Wang. Frido: Feature pyramid diffusion for complex scene image synthesis. *arXiv preprint arXiv:2208.13753*, 2022. [3](#)
- [19] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2334–2343, 2017. [6](#)
- [20] Runyang Feng, Yixing Gao, Xueqing Ma, Tze Ho Elden Tse, and Hyung Jin Chang. Mutual information-based temporal difference learning for human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17131–17141, June 2023. [1](#)
- [21] Yixing Gao, Hyung Jin Chang, and Yiannis Demiris. User modelling for personalised dressing assistance by humanoid robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1840–1845. IEEE, 2015. [1](#)
- [22] Yixing Gao, Hyung Jin Chang, and Yiannis Demiris. Iterative path optimisation for personalised dressing assistance using vision and force information. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4398–4403. IEEE, 2016. [1](#)
- [23] Zigang Geng, Ke Sun, Bin Xiao, Zhaoxiang Zhang, and Jingdong Wang. Bottom-up human pose estimation via disentangled keypoint regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14676–14686, 2021. [2](#)
- [24] Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, and Du Tran. Detect-and-track: Efficient pose estimation in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 350–359, 2018. [6](#)
- [25] Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022. [3](#)
- [26] Zhangxuan Gu, Haoxing Chen, Zhuoer Xu, Jun Lan, Changhua Meng, and Weiqiang Wang. Diffusioninst: Diffusion model for instance segmentation. *arXiv preprint arXiv:2212.02773*, 2022. [2](#), [3](#)
- [27] Hengkai Guo, Tang Tang, Guozhong Luo, Riwei Chen, Yongchen Lu, and Linfu Wen. Multi-domain pose net-

- work for multi-person pose estimation and tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 6
- [28] Kota Hara and Rama Chellappa. Computationally efficient regression on a dependency graph for human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3390–3397, 2013. 2
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [30] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 1, 2, 3, 5
- [31] Rongjie Huang, Zhou Zhao, Huadai Liu, Jinglin Liu, Chenye Cui, and Yi Ren. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2595–2605, 2022. 3
- [32] Jihye Hwang, Jieun Lee, Sungheon Park, and Nojun Kwak. Pose estimator and tracker using temporal flow maps for limbs. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019. 6
- [33] Umar Iqbal, Anton Milan, and Juergen Gall. Posetrack: Joint multi-person pose estimation and tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2, 6
- [34] Sheng Jin, Wentao Liu, Wanli Ouyang, and Chen Qian. Multi-person articulated tracking with spatial and temporal embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5664–5673, 2019. 6
- [35] Martin Kiefel and Peter Vincent Gehler. Human pose estimation with fields of parts. In *European conference on computer vision*, pages 331–346. Springer, 2014. 2
- [36] Jiefeng Li, Siyuan Bian, Ailing Zeng, Can Wang, Bo Pang, Wentao Liu, and Cewu Lu. Human pose regression with residual log-likelihood estimation. *arXiv preprint arXiv:2107.11291*, 2021. 1
- [37] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *arXiv preprint arXiv:2205.14217*, 2022. 3
- [38] Yanjie Li, Shoukui Zhang, Zhicheng Wang, Sen Yang, Wankou Yang, Shu-Tao Xia, and Erjin Zhou. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11313–11322, 2021. 1, 2
- [39] Zhenguang Liu, Haoming Chen, Runyang Feng, Shuang Wu, Shouling Ji, Bailin Yang, and Xun Wang. Deep dual consecutive network for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 525–534, 2021. 1, 2, 6
- [40] Zhenguang Liu, Runyang Feng, Haoming Chen, Shuang Wu, Yixing Gao, Yunjun Gao, and Xiang Wang. Temporal feature alignment and mutual information maximization for video-based human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11016, 2022. 1, 2, 6, 7, 8
- [41] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 3
- [42] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. 1
- [43] Yue Luo, Jimmy Ren, Zhouxia Wang, Wenxiu Sun, Jinshan Pan, Jianbo Liu, Jiahao Pang, and Liang Lin. Lstm pose machines. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5207–5215, 2018. 5
- [44] Naoki Matsunaga, Masato Ishii, Akio Hayakawa, Kenji Suzuki, and Takuya Narihira. Fine-grained image editing by pixel-wise guidance using diffusion models. *arXiv preprint arXiv:2212.02024*, 2022. 3
- [45] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 2, 5
- [46] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3
- [47] Tomas Pfister, James Charles, and Andrew Zisserman. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1913–1921, 2015. 2
- [48] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-tts: A diffusion probabilistic model for text-to-speech. In *International Conference on Machine Learning*, pages 8599–8608. PMLR, 2021. 3
- [49] Yaadhav Raaj, Haroon Idrees, Gines Hidalgo, and Yaser Sheikh. Efficient online multi-person 2d pose tracking with recurrent spatio-temporal affinity fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4620–4628, 2019. 6
- [50] Umer Rafi, Andreas Doering, Bastian Leibe, and Juergen Gall. Self-supervised keypoint correspondences for multi-person pose estimation and tracking in videos. In *European Conference on Computer Vision*, pages 36–52. Springer, 2020. 6
- [51] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019. 6
- [52] Mengwei Ren, Mauricio Delbracio, Hossein Talebi, Guido Gerig, and Peyman Milanfar. Image deblurring with domain generalizable diffusion models. *arXiv preprint arXiv:2212.01789*, 2022. 1
- [53] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted*

- Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 2
- [54] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 3
- [55] Benjamin Sapp, Alexander Toshev, and Ben Taskar. Cascaded models for articulated pose estimation. In *European conference on computer vision*, pages 406–420. Springer, 2010. 1
- [56] Luca Schmidtke, Athanasios Vlontzos, Simon Ellershaw, Anna Lukens, Tomoki Arichi, and Bernhard Kainz. Unsupervised human pose estimation through transforming shape templates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2484–2494, 2021. 1
- [57] Dahu Shi, Xing Wei, Liangqi Li, Ye Ren, and Wenming Tan. End-to-end multi-person pose estimation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11069–11078, 2022. 1
- [58] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 1, 3
- [59] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3, 5
- [60] Jie Song, Limin Wang, Luc Van Gool, and Otmar Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4220–4229, 2017. 2
- [61] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 3
- [62] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 3
- [63] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 1, 2, 6, 8
- [64] Tze Ho Elden Tse, Daniele De Martini, and Letizia Marchegiani. No need to scream: Robust sound-based speaker localisation in challenging scenarios. In *Social Robotics: 11th International Conference, ICSR 2019*, pages 176–185. Springer, 2019. 1
- [65] Tze Ho Elden Tse, Kwang In Kim, Ales Leonardis, and Hyung Jin Chang. Collaborative learning for hand and object reconstruction with attention-guided graph convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1664–1674, 2022. 1
- [66] Tze Ho Elden Tse, Zhongqun Zhang, Kwang In Kim, Ales Leonardis, Feng Zheng, and Hyung Jin Chang. S²Contact: Graph-based network for 3D hand-object contact estimation with semi-supervised learning. In *European Conference on Computer Vision*, pages 568–584. Springer, 2022. 1
- [67] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2
- [68] Dongkai Wang and Shiliang Zhang. Contextual instance decoupling for robust multi-person pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11060–11068, 2022. 1
- [69] Fang Wang and Yi Li. Beyond physical connections: Tree models in human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 596–603, 2013. 1
- [70] Manchen Wang, Joseph Tighe, and Davide Modolo. Combining detection and tracking for human pose estimation in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11088–11096, 2020. 2, 6
- [71] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016. 2, 5
- [72] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [73] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. 2, 6
- [74] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose flow: Efficient online pose tracking. *arXiv preprint arXiv:1802.00977*, 2018. 6
- [75] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vit-pose: Simple vision transformer baselines for human pose estimation. *arXiv preprint arXiv:2204.12484*, 2022. 1, 2, 3
- [76] Dongchao Yang, Jianwei Yu, Helin Wang, Wen Wang, Chao Weng, Yuexian Zou, and Dong Yu. Diffsound: Discrete diffusion model for text-to-sound generation. *arXiv preprint arXiv:2207.09983*, 2022. 3
- [77] Sen Yang, Zhibin Quan, Mu Nie, and Wankou Yang. Transpose: Keypoint localization via transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11802–11812, 2021. 2
- [78] Yiding Yang, Zhou Ren, Haoxiang Li, Chunlun Zhou, Xinchao Wang, and Gang Hua. Learning dynamics via graph neural networks for human pose estimation and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8074–8084, 2021. 6
- [79] Dongdong Yu, Kai Su, Jia Sun, and Changhu Wang. Multi-person pose estimation for pose tracking with enhanced cascaded pyramid network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 6

- [80] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution vision transformer for dense predict. *Advances in Neural Information Processing Systems*, 34:7281–7293, 2021. [1](#), [2](#)
- [81] Jiabin Zhang, Zheng Zhu, Wei Zou, Peng Li, Yanwei Li, Hu Su, and Guan Huang. Fastpose: Towards real-time pose estimation and tracking via scale-normalized multi-task networks. *arXiv preprint arXiv:1908.05593*, 2019. [6](#)
- [82] Xiaoqin Zhang, Changcheng Li, Xiaofeng Tong, Weiming Hu, Steve Maybank, and Yimin Zhang. Efficient human pose estimation via parsing a tree structure based human model. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1349–1356. IEEE, 2009. [1](#)