Improved 3D Markerless Mouse Pose Estimation Using Temporal Semi-supervision

Tianqing Li[1] · Kyle S. Severson[2] · Fan Wang[2] · Timothy W. Dunn[1]

Abstract

Three-dimensional markerless pose estimation from multi-view video is emerging as an exciting method for quantifying the behavior of freely moving animals. Nevertheless, scientifically precise 3D animal pose estimation remains challenging, primarily due to a lack of large training and benchmark datasets and the immaturity of algorithms tailored to the demands of animal experiments and body plans. Existing techniques employ fully supervised convolutional neural networks (CNNs) trained to predict body keypoints in individual video frames, but this demands a large collection of labeled training samples to achieve desirable 3D tracking performance. Here, we introduce a semi-supervised learning strategy that incorporates unlabeled video frames via a simple temporal constraint applied during training. In freely moving mice, our new approach improves the current state-of-the-art performance of multi-view volumetric 3D pose estimation and further enhances the temporal stability and skeletal consistency of 3D tracking.

1 Introduction

In 3D pose estimation, the positions of user-defined body keypoints are inferred from images to reconstruct body kinematics (Desmarais et al., 2021). Precise pose measurement is a long-standing computer vision research problem with a myriad of applications, including to human-computer interfaces, autonomous driving, virtual and artificial reality, and robotics (Sarafianos et al., 2016). Specialized hardware and deep learning empowered algorithmic advances have inspired new developments in the field, with the ultimate goal

✉ Timothy W. Dunn
timothy.dunn@duke.edu

Tianqing Li
tianqing.li@duke.edu

Kyle S. Severson
kyleseve@mit.edu

Fan Wang
fan_wang@mit.edu

[1] Department of Biomedical Engineering, Pratt School of Engineering, Duke University, Durham 27708, NC, USA

[2] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge 02140, MA, USA

to recover 3D body poses in natural, occlusive environments in real time. While most research and development have thus far focused on human body tracking, there has been a growing push in the biological research community to extend 3D human pose estimation techniques to animals. Precise quantification of animal movement is critical for understanding the neural basis of complex behaviors and neurological diseases (Marshall et al., 2022). The latest generation of tools for animal behavior quantification ditch traditionally coarse and ad hoc measurements for 2D and 3D pose estimation with convolutional neural networks (CNNs) (Pereira et al., 2019; Mathis et al., 2018; Pereira et al., 2022; Günel et al., 2019; Bala et al., 2020; Gosztolai et al., 2021; Dunn et al., 2021).

Nevertheless, the majority of state-of-the-art 3D animal pose estimation techniques are fully supervised, and their performance depends on large collections of 2D and 3D annotated training samples. Large-scale, well-curated animal 3D pose datasets are still rare, making it difficult to achieve consistent results on real-world data captured under varying experimental conditions. Marker-based motion capture techniques (Mimica et al., 2018; Marshall et al., 2021) enable harvesting of precise and diverse 3D body pose measurements, but they are difficult to deploy in freely moving animals and can potentially perturb natural behaviors. Man-

ual annotation of animal poses therefore often becomes mandatory. However, manual annotation is time-consuming, and it can become difficult for human annotators to precisely localize body landmark positions under nonideal lighting conditions or heavy (self-) occlusion of the body. Although the influence of label noise has not yet been closely examined for pose estimation, overfitting to these inherently ambiguous labels might adversely affect model performance, as it does in image classification (Patrini et al., 2017). In addition to issues with data scarcity, fully supervised training schemes are often limited by the quality of training data. Even when using hundreds of training samples, the performance of fully supervised 3D pose estimation models can be inconsistent (Wu et al., 2020), especially when deployed in new environments and subjects.

This label scarcity has become a major bottleneck in the current animal 3D pose estimation workflows, limiting model performance, generalization to different environments and species, and comprehensive performance analysis. In recent years, the success of semi-supervised (Berthelot et al., 2019) and unsupervised deep learning (He et al., 2020; Chen et al., 2020) methodologies has presented new possibilities for mitigating annotation burden. Rather than relying solely on task-relevant information provided by human supervision, these approaches exploit the abundant transferable features embedded in unlabeled data, resulting in robustness to annotation deprivation and better generalization capacity.

In this paper, we introduce a semi-supervised framework which seamlessly integrates with the current state-of-the-art 3D rodent pose estimation approach (Dunn et al., 2021) to enhance tracking performance in low annotation regimes. The core of our approach is additional regularization of body landmark localization using a Laplacian temporal prior 1. This encourages smoothness in 3D tracking trajectories without imposing hard constraints, while expanding supervisory signals to include both human-annotated labels and the implicit cues abundant in unlabeled video data. To further reduce reliance on large labeled datasets, we also emphasize a new set of evaluation protocols that operate on unlabeled frames, thus providing more comprehensive performance assessments for markerless 3D animal pose estimation algorithms. We have collected and validated our proposed method on a new multi-view video-based mouse behavior dataset with 2D and 3D pose annotations, which have released to the community. Compared to state-of-the-art approaches in both animal and human pose estimation, our method improves keypoint localization accuracy by 15 to 60% in low annotation regimes, achieves better tracking stability and anatomical consistency, and is qualitatively more robust during identified difficult poses.

Our main contributions can be summarized as follows:

(1) We introduce a new state-of-the-art approach by leveraging temporal supervision in 3D mouse pose estimation.
(2) We release a new multi-view 3D mouse pose dataset consisting of freely moving, naturalistic behaviors to the community.
(3) We benchmark the performance of a broad range of contemporary pose estimation algorithms using the new dataset.
(4) We designate a comprehensive set of evaluation metrics for performance assessment of animal pose estimation approaches.

2 Related Work

2.1 3D Animal Pose Estimation

There are currently three primary categories of 3D animal pose estimation techniques. The first category encompasses multi-view approaches based on triangulation of 2D keypoint estimates (Mathis et al., 2018; Günel et al., 2019; Bala et al., 2020; Karashchuk et al., 2021). These are typically lightweight in terms of model training and inference and are improved by post hoc spatial-temporal filtering (Karashchuk et al., 2021) when measuring freely moving behavior, where occlusions are ubiquitous. The second category leverages multi-view geometric information during end-to-end training. Zimmermann et al. (2020) and Dunn et al. (2021) use 3D CNNs to process volumetric image representations obtained via projective geometry, whereas (Yao et al., 2019) propose a self-supervised training scheme based on cross-view epipolar information. These techniques improve 3D tracking accuracy and consistency by exploiting multi-view features during training, although they are more computationally demanding. The third category comprises learned transformations of monocular 2D pose estimates into 3D space (Gosztolai et al., 2021; Bolaños et al., 2021). Monocular 3D pose estimation is an exciting and important advance in flexibility, but unavoidable 3D ambiguities currently limit its performance compared to multi-view techniques (Iskakov et al., 2019; Bolaños et al., 2021).

Despite the recent acceleration in method development, it remains challenging to build 3D animal pose estimation algorithms that achieve scientifically precise performance flexibly across diverse environments and species. Compared to humans, lab animals such as mice and rats are much smaller in scale, less articulated, and bear higher appearance similarities among different individuals (Moskvyak et al., 2020), which limits the availability of discriminable features for body part tracking and annotation. Because of the drastic differences in animal body profiles across species, (e.g. cheetahs vs. flies), it is also difficult to leverage the universal skeleton models and large-scale pretraining datasets that power the
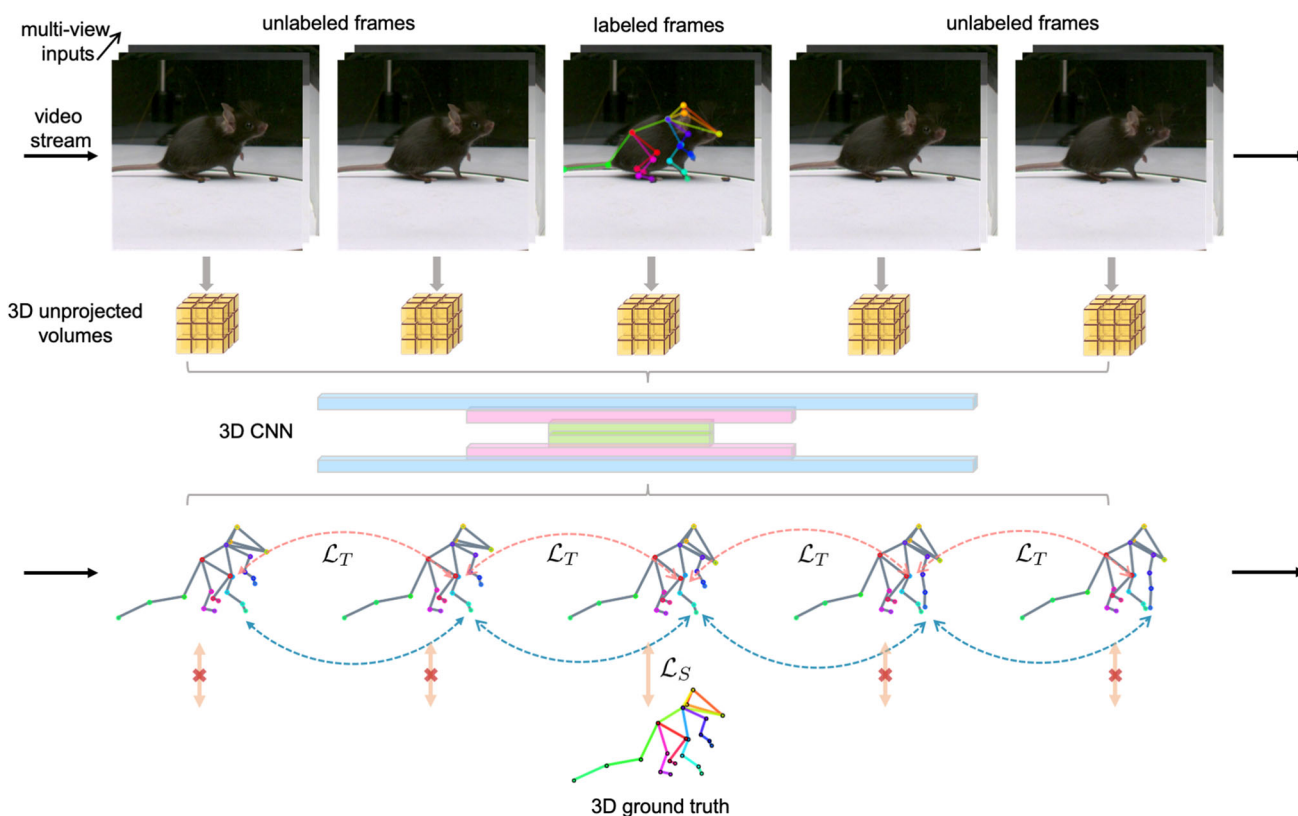
Fig. 1 Method overview. Our multi-view volumetric approach constructs a 3D image feature grid using projective geometry for each timepoint in videos. A 3D CNN (UNet) processes batches of temporally contiguous volumetric inputs and directly predicts 3D keypoint positions. We then combine a traditional supervised regression loss with an unsupervised temporal consistency loss for training. While the regression loss $\mathcal{L}_S$ is applied only on labeled video frames, which are sparsely distributed across video recordings, the unsupervised temporal loss $\mathcal{L}_T$ operates over both labeled and unlabeled frames

impressive tracking performance in humans (Cao et al., 2019; Wu et al., 2020). It is imperative that we develop algorithms that more efficiently use the limited resources available for animals.

2.2 Semi-supervised and Unsupervised Pose Estimation

Semi-supervised and unsupervised learning schemes reduce the reliance on laborious data annotation currently bottlenecking large-scale supervised training. These schemes learn from the implicit structure and distribution of unlabeled data and can utilize knowledge of universal principles, such as physics and geometry, to improve tracking performance.

Inspired by classic multi-view stereo 3D reconstruction, many works in 3D human pose estimation utilize annotation-free geometric supervision in the form of multi-view consistency (Rhodin et al., 2018; Kocabas et al., 2019; Iqbal et al., 2020; Wandt et al., 2021), 3D-to-2D reprojection consistency (Wandt & Rosenhahn, 2019; Chen et al., 2019), and geometry-aware 3D representation learning (Rhodin et

al., 2018). Training constraints with respect to consistent bone length, valid ranges of joint angles, and body symmetry (Wu et al., 2020; Spurr et al., 2020; Dabral et al., 2018; Pavllo et al., 2019) can also encourage biomechanically-plausible tracking results. Exploiting temporal context is also effective, as we discuss in the next section. Appropriate use of these implicit supervision signals results in consistent and robust pose estimates using only a small fraction of the labeled data required for fully supervised approaches.

2.3 Temporal 3D Pose Estimation

The temporal nature of behavior provides information that can be harnessed to improve 3D pose estimation. Intuitively, movement progresses continuously through time in 3D space, providing a strong prior for future poses given their temporal history—body movement trajectories evolve smoothly and are bounded by plausible, physiological velocities. The spatial displacement between consecutive poses should therefore be small, exhibiting relative consistency or smoothness along the time dimension. Pose estimates from

static, temporally isolated observations ignore these intuitive constraints.

Previous 3D pose estimation algorithms have incorporated temporal information in several different ways. Given a sequence of pose predictions, temporal consistency can be introduced as part of the post-processing optimization that refines initial 2D (prior to triangulation) or 3D keypoint estimates (Bala et al., 2020; Joska et al., 2021; Karashchuk et al., 2021; Zhang et al., 2021). Temporal consistency assumptions have also been used for filtering out invalid pseudolabels used for self-supervision (Mu et al., 2020).

Another popular scheme for exploiting temporal information for 3D pose estimation is to build models that infer pose from spatiotemporal inputs, using either recurrent neural networks (Hossain et al., 2018), temporal CNNs (Pavllo et al., 2019), or spatial-temporal graphical models (Wang et al., 2020). Hossain and Little (Hossain et al., 2018) processed 2D pose sequences using layer-normalized LSTMs to produce temporally consistent 3D poses. Other works have used temporal CNNs for similar purposes (Pavllo et al., 2019; Liu et al., 2020). Temporal information can also be explicitly encoded and appended to model input using apparent motion estimations such as optical flow (Liu et al., 2021).

Other approaches incorporate temporal information as a form of regularization during training. By employing a temporal smoothness constraint, one enforces the assumption that joint positions should not displace significantly over short periods of time (Wu et al., 2020; Wang et al., 2020), encouraging learned temporal consistency in pose predictions. Critically, these temporal constraints can be applied to unlabeled video frames, providing an avenue for semi- and unsupervised learning. Chen et al. (2021) further exploited temporal consistency in hand pose estimates along both forward and backward video streaming directions to establish an effective self-supervised learning scheme. Our approach is most similar to Wu et al. (2020), in that we incorporate a temporal smoothness constraint in the learning objective to support a semi-supervised scheme. But we employ this constraint with multi-view, volumetric 3D pose estimation during freely moving, naturalistic behavior, rather than during monocular 2D pose estimation in restrained animals.

2.4 Pose Evaluation Metrics

In this manuscript we also report a complementary set of performance metrics that provides more comprehensive benchmarks for sparsely labeled 3D animal pose data. The cornerstone metrics of the field are Euclidean distance errors relative to ground-truth 3D keypoints: mean per-joint position error (MPJPE), and, sometimes, PA-MPJPE, which evaluates MPJPE after rigid alignment of 3D predictions to ground-truth poses. Although these evaluation protocols convey an imperative assessment of a model's landmark localization capability, they fall short for most markerless animal pose datasets, where 3D keypoint ground-truth is derived from noisy manual labeling only in a small subset of video frames.

Unlike in large-scale human benchmarks, in animals these position error metrics do not reflect the large extant diversity of possible poses and are prone to overestimating performance. Human3.6M (Ionescu et al., 2013) and HumanEva (Sigal et al., 2010) employ motion capture systems to acquire comprehensive ground-truth labels over hundreds of thousands of frames, spanning multiple human actors and dozens of action categories. Similar evaluation is nearly impossible for most markerless 3D animal pose datasets, where acquisition of 3D labels requires laborious human annotation.

Single-frame position errors over sparsely labeled recordings also ignore whether models capture the continuous and smooth nature of movement. Models with the same mean position error on a small subset of samples can diverge significantly, and pathologically, in unlabeled frames. We illustrate this in Fig. 2a, which shows a set of synthetic movement trajectories. The three noisy traces all have the same average position relative to 100 points sampled evenly from the ground truth, yet the traces represent distinct, and erroneous, movement patterns. The issue can become even more pronounced when ground-truth labels are sparse and unevenly distributed, as is the case with most animal datasets. The fidelity of predictions on unlabeled data could be captured using temporal metrics. To quantify the temporal consistency of predictions (Pavllo et al., 2019). However, works in animal pose estimation do not typically incorporate quantitative temporal metrics on unlabeled frames, although some have presented qualitative comparisons to keypoint position (Wu et al., 2020) or movement velocity traces (Karashchuk et al., 2021) or reported quantitative errors within manually labeled frames (Karashchuk et al., 2021).

Finally, manually annotated 3D pose ground-truth is inherently noisy and exhibits substantial intra- and inter-labeler variability. We analyzed the coefficients of variation ($CV = \frac{\sigma}{\mu}$) (Reed et al., 2002), which measures the degree of data dispersion relative to its mean, for the lengths of 22 body segments connecting keypoints in our manually labeled mouse dataset (details in Section 4.1). Although the keypoints are intended to represent body joints, between which the lengths of body segments should remain constant, independent of pose, we found a 10% to 20% deviation in length for the majority of segments (Fig. 2b). This aleatoric uncertainty in the ground-truth labels will propagate to position errors.

Given these issues, we argue that it is important to establish more diverse evaluation protocols for markerless 3D animal pose estimation. These protocols should ideally capture temporal and anatomical variances in both labeled and unlabeled frames. In addition to our new semi-supervised training scheme, we introduce two new consistency metrics
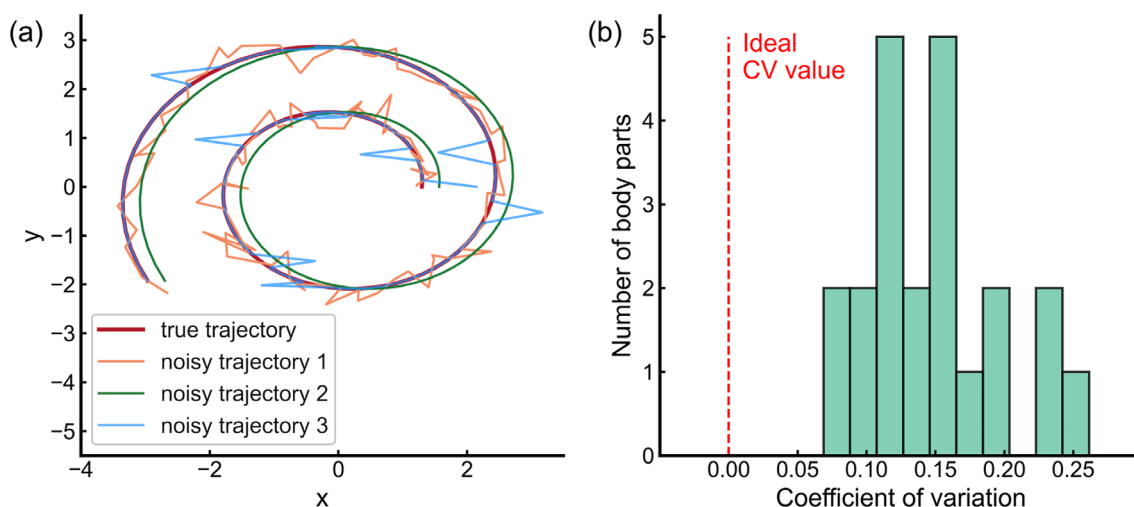
Fig. 2 a Ambiguity in absolute position error analysis. In this simulated example, we present three noisy trajectories with the same absolute point position errors with respect the true spiral trajectory. **b** Histogram of body segment length variation in manually labeled mouse data. We compute the coefficient of variation (CV) for the lengths of 22 body segments. While CV values should ideally be close to 0, we instead observed notable amounts of length variation in all body segments. This illustrates the noise present in manually labeled 3D poses

that resolve differences between models not captured by standard position errors, and these new metrics do not rely on large numbers of ground-truth annotations.

2.5 3D Animal Pose Datasets and Benchmarks

Despite the critical importance of large-scale, high-quality datasets for developing 3D animal pose estimation algorithms (Jain et al., 2020), such resources are relatively uncommon compared to what is available for 3D human pose. Animal datasets are not easily applied across species, due to differences in body plans, and high-throughput marker-based motion capture techniques are challenging to implement in freely-moving, small-sized animals. Nevertheless, multiple 3D animal pose datasets have been released in recent years, including in dogs (Kearney et al., 2020), cheetahs (Joska et al., 2021), rats (Dunn et al., 2021; Marshall et al., 2021), flies (Günel et al., 2019), and monkeys (Bala et al., 2020). But in mice, by far the most commonly used mammalian model organism in biomedical research (Ellenbroek & Youn, 2016), large-scale pose datasets are still lacking. The LocoMouse dataset (Machado et al., 2015) contains annotated 3D keypoints in animals walking down a linear track. While being a valuable resource for developing gait tracking algorithms, the dataset does not represent the diversity of mouse poses composing the naturalistic behavioral repertoire. Several 3D mouse datasets also accompany published manuscripts (Zimmermann et al., 2020), but they are limited in the number of total annotated frames. Here we provide a new, much larger 3D mouse pose dataset consisting of 6.7 million frames with 310 annotated 3D poses (1860 annotated frames in 2D) on 5 mice engaging in freely moving, naturalistic behaviors, which we make publicly available as a resource for the community. We also utilize the scale of our dataset to benchmark a collection of popular 3D pose estimation algorithms and assess the impact of temporal constraints on performance, providing guidance on the development of suitable strategies for quantifying mouse behavior in three dimensions.

3 Methods

3.1 Volumetric Representation

Following recent computer stereo vision methods (Kar et al., 2017; Iskakov et al., 2019; Zimmermann et al., 2020; Dunn et al., 2021), we construct a geometrically-aligned volumetric input V_t from multi-view video frames at each timepoint t and estimate 3D pose from them using a 3D CNN.

As memory limitations restrict the size of the 3D volume ($64 \times 64 \times 64$ voxels in our case), to increase its spatial resolution, we center the volume at the inferred 3D centroid of the animal. This centroid is inferred by triangulating 2D centroids detected in each camera view using a standard 2D UNet (Ronneberger et al., 2015), except with half the number of channels in each convolutional layer. For triangulation, we take the median of all pairwise triangulations across views. We then create an axis-aligned 3D grid cube centered at the 3D centroid position, which bounds the animal in 3D world space. We use $N = 64$ voxels per grid cube side, resulting in an isometric spatial resolution of 1.875 mm per voxel.

Here, we briefly review the volume generation process. After initialization, 3D grids are populated with 2D image RGB pixel values from each camera using projective geometry. With known camera extrinsic (rotation matrix R, translation vector t) and intrinsic parameters K, a 2D image $\mathcal{F}$ can be unprojected along the viewing rays as they intersect with the 3D grid. In practice, rather than performing actual ray tracing, the center coordinates of each 3D voxel $X_{i,j,k}$ is projected onto the target 2D image plane by $K[R \mid t]X_{i,j,k}$ and the value of $X_{i,j,k}$ is set by bilinear sampling from the image at the projected point (Kar et al., 2017). The unprojected image volumes from different views are concatenated along the channel dimension, resulting in a $N \times N \times N \times (N_{cam} * C)$-sized volumetric input, where C is the channel dimension size of each input view ($C = 3$ for RGB images). While we sample directly from 3-channel RGB images to reduce memory footprint and computation costs, other approaches unproject features extracted by 2D CNNs (Iskakov et al., 2019; Tu et al., 2020; Zimmermann et al., 2020).

The unprojected image volumes are then processed by a 3D UNet (implementation details in Sect. 4.5), producing volumetric heatmaps associated with different keypoints. The differentiable expectation operation *soft argmax* (Nibali et al., 2018; Sun et al., 2018) is applied along spatial axes to infer the numerical coordinates of each keypoint.

3.2 Unsupervised Temporal Loss

At high frame rates, the per-frame velocity of animals is low and their overall movement trajectory should typically be smooth. We encode these assumptions as an unsupervised temporal smoothness loss $\mathcal{L}_T(\cdot)$ that can be easily integrated with heatmap-based pose estimation approaches.

Consider the inputs to the network to be a set of temporally consecutive chunks $\mathcal{T}$ where each chunk $\mathcal{T}_n$ consists of 3D volumetric representations constructed from c adjacent timepoints $\mathcal{T}_n = \{V_{t_i}, \ldots, V_{t_{i+c}-1}\}$, where c specifies the time span covered by the unsupervised loss.

Given the 3D keypoint coordinates predicted by the 3D CNN $\{J_{t,j} \mid t_i \le t \le t_{i+c-1}, 1 \le j \le N_J\}$ from one temporal chunk $\mathcal{T}_n$, the temporal smoothness loss penalizes the keypoint-wise position divergence across consecutive frames, which is equivalent to constraining the movement velocity within the temporal window.

$$\mathcal{L}_T(\{J_{t,j}\}) = \frac{1}{c} \sum_{t=t_i}^{t_{i+c-1}} \frac{1}{N_J} \sum_{j=1}^{N_J} d(J_{t,j}, J_{t+1,j}) \tag{1}$$

where N_J is the number of 3D keypoints and d is the distance metric used for comparing displacement across timepoints.

This general formulation does not enforce limitations on the choice of distance metric, but empirically we found that L1 distance performed better than L2-norm Euclidean distance. Though it is difficult to give a theoretical explanation for this observation, the underlying reason could be similar to that for L1 total variation regularization in optical flow estimation. Formulating the smoothness constraint as a Laplacian prior allows discontinuity in the motion and is well known to be more robust to data outliers compared to quadratic regularizers (Wedel et al., 2009). We have therefore used an L1 distance metric for all experiments presented in the later sections.

3.3 Supervised Pose Regression Loss

The unsupervised temporal loss on its own is insufficient and will result in mode degeneracy where the network learns to produce identical poses for all input samples. We therefore also include a standard supervised pose regression loss over a small set of labeled frames during training. Given the ground-truth and predicted 3D keypoint coordinates J_t and $\hat{J}_t$, the supervised regression loss is defined as

$$\mathcal{L}_S(J_t, \hat{J}_t) = \frac{1}{N_J} \sum_{j=1}^{N_J} d(J_{t,j}, \hat{J}_{t,j}). \tag{2}$$

We use L1 distance for computing the joint distances over L2 distance metric based on empirical results, which agrees with the results of Sun et al. on 3D human pose estimation (Sun et al., 2018).

4 Experiments

4.1 Dataset

For performance evaluation, we collected a total of five 1152×1024 pixels color video recordings from 6 synchronized cameras surrounding a cylindrical arena. We direct the reader to "Appendix B" and Supplementary Video 1 for more details on the 3D mouse pose dataset. Each set of recordings corresponds to a different mouse (M1, M2, M3, M4, M5). M1 and M2 were recorded for 3 minutes and M3, M4, M5 were recorded for 60 minutes. The number of manually annotated 3D ground-truth timepoints for 22 body keypoints is n = 81, 91, 48, 44 and 46 from each recording, respectively (486, 546, 288, 264, and 276 total annotated video frames). Out of the 22 keypoints, 3, 4, 6, 6, and 3 are located on the animal's head, trunk, forelimbs, hindlimbs, and tail, respectively. Notice that the two keypoints at the middle and end of the tail were excluded from quantitative evaluations presented in this paper, as they were often cropped outside the

bounds of the 3D grids. This results in a total of 20 body keypoints and 22 corresponding body segments used for analysis.

We allocated n = 172 from M1 and M2 for training and n = 48 from M3 for internal validation. We report all metrics using data from M4 and M5 (n = 90 labeled timepoints, plus unlabeled timepoints for additional temporal and anatomical consistency metrics), which were completely held out from training or model selection. We also simulated low annotation conditions by randomly selecting 5% (n = 8), 10% (n = 17) and 50% (n = 86) from the training samples and compared with the full annotation 100% condition.

4.2 Evaluation Metrics

4.2.1 Localization Accuracy

We adopt the three common protocols used in 3D human pose estimation for evaluating the landmark localization accuracy of different models. Metric results are averaged over all the labeled timepoints as described in Sect. 4.1.

- Protocol #1: Mean per-joint position error (**MPJPE**) evaluates the mean joint-wise 3D Euclidean distances between the prediction and ground truth keypoint positions. For J keypoints in a single frame,

$$\text{MPJPE}(\mathbf{J}) = \frac{1}{N_J} \sum_{j=1}^{N_J} \|\hat{\mathbf{J}}_j - \mathbf{J}_j^{gt}\|_2$$

- Protocol #2: Procrustes Analysis MPJPE (**PA-MPJPE**) reports the MPJPE values after rigidly aligning the landmark predictions (translation and rotation) with the ground-truth.
- Protocol #3: Normalized MPJPE (**N-MPJPE**) assesses the scale-insensitive MPJPE estimation errors by respectively normalizing the prediction and ground-truth landmarks by their norm (Rhodin et al., 2018).

4.2.2 Temporal Smoothness

The aforementioned single-frame evaluation metrics are inadequate for capturing the importance of temporal smoothness in videos. We therefore also report a mean per-joint temporal deviation (**MPJTD**) metric, which we define simply as the mean absolute value of first-order derivative of predicted pose sequences. We used $T = 10000$ continuous frames from recordings of mouse M5 for this evaluation.

$$\text{MPJTD}(\mathbf{J}) = \frac{1}{T-1} \frac{1}{\cdot N_J} \sum_{j=1}^{N_J} \sum_{t=1}^{T-1} |\mathbf{J}_{t,j} - \mathbf{J}_{t+1,j}|$$

4.2.3 Body Skeleton Consistency

Although not explicitly constrained during training, the anatomical consistency of predictions is an important component of model tracking performance. Inspired by the analysis of (Karashchuk et al., 2021), we examined the mean and standard deviation of the estimated length of 22 body segments over 10000 continuous frames from M4 for this analysis.

4.3 Training Strategies

To evaluate the influence of temporal training, we designed four different model training schemes that were each applied to the 5%, 10%, 50% and 100% annotation conditions.

Baseline/DANNCE (Dunn et al., 2021) We employ the multi-view volumetric method presented by Dunn et al. as the baseline comparison. All baseline models are trained solely with the supervised regression L1 loss over the labeled frames.

Baseline + smoothing No changes are made during the training; instead, the predictions from the baseline models are smoothed in time for each keypoint, with a set of different smoothing strategies.

Temporal baseline During training, each batch contains exactly one labeled sample with three additional unlabeled samples drawn from its local neighborhood. This scheme ensures a balance between supervised and unsupervised loss throughout the optimization. The models were then jointly trained with $\mathcal{L}_S$ and $\mathcal{L}_T$.

Temporal + extra In addition to the partially labeled training batches used in temporal baseline model training, the training set contains N_u completely unlabeled, temporally consecutive chunks included only in the unsupervised temporal loss.

For experiments conducted under lower annotation conditions, 5%, 10% and 50%, we use respectively 95% ($N_u = 163$), 90% ($N_u = 154$) and 50% ($N_u = 86$) unlabeled chunks with respect to the entire training set. This aimed to match the number of samples used in the 100% baseline and temporal baseline models. For experiments using 100% of the training data, we add 20% ($N_u = 34$) extra unlabeled temporal chunks.

4.4 Comparison with State-of-the-Art Approaches

We compare the performance of our proposed approach against other contemporary animal and human pose estimation methods. Specifically, we have replicated and evaluated the following approaches on the mouse dataset:

2D animal pose estimation DeepLabCut (DLC) (Mathis et al., 2018) is a widely adopted toolbox for markerless pose
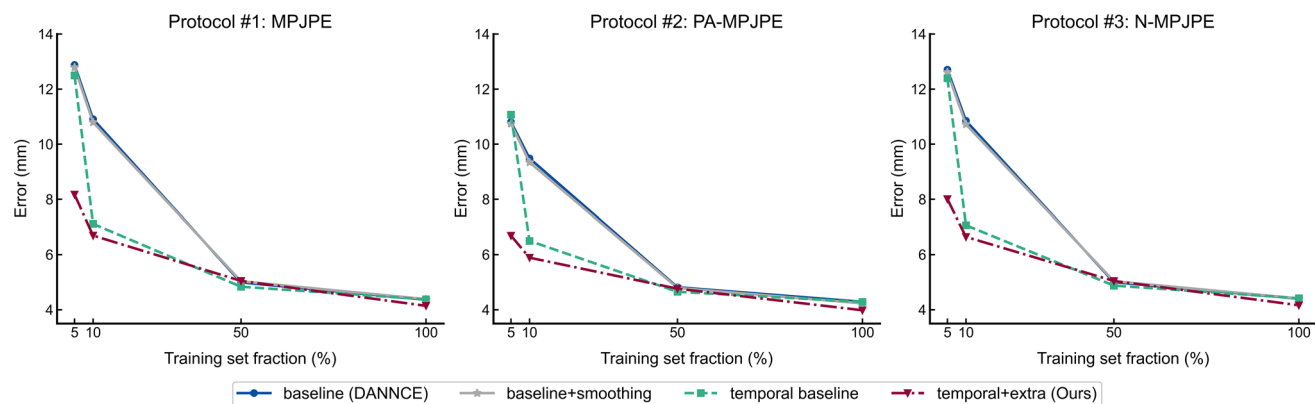
Fig. 3 Qualitative comparison of landmark localization performance over different annotation conditions. We randomly selected 5% (n = 8), 10% (n = 17) and 50% (n = 85) of the training set to simulate low annotation regimes. Temporal supervision generally improved performance on all three localization protocols compared to the baseline models, especially with limited access to the training data. Similar improvement cannot be achieved via post hoc smoothing of the predicted movement trajectories

Table 1 Quantitative comparison with other state-of-the-art 2D and 3D animal and human pose estimation methods

Protocol 1 (3D MPJPE, mm)				
	Training set fraction			
	5%	10%	50%	100%
2D pose estimation methods (+ post hoc triangulation)				
DLC[†] (Mathis et al., 2018)	11.09	11.05	9.89	8.90
SimpleBaseline[†] (Xiao et al., 2018)	18.09	14.61	7.36	5.95
SimpleBaseline	18.56	16.58	8.35	6.69
DLC + *soft argmax*	11.03	9.22	6.35	6.47
DLC + 2D variant of our temporal constraint*	8.54	9.12	5.95	6.03
3D monocular pose estimation methods				
Temporal Convolution* (Pavllo et al., 2019)	–	–	–	17.63
3D multi-view pose estimation methods				
Learnable Triangulation[†] (Iskakov et al., 2019)	18.77	15.66	8.97	6.31
DANNCE (Dunn et al., 2021)	12.87	10.90	4.99	4.36
Ours (temporal baseline)*	12.49	7.11	4.83	4.37
Ours (temporal + extra)*	8.17	6.69	5.04	4.14

We report the absolute 3D MPJPE in millimeters for each approach using four different fractions of training data. The methods that use ground truth 2D bounding boxes during inference are masked by †. The methods that use temporal information during training are masked by ∗. For the monocular approach, the reported metric results are separately computed and averaged across all camera views

estimation of animals, which expanded on the previous state-of-the-art method DeeperCut (Insafutdinov et al., 2016). We followed the default architecture and training configurations using ResNet-50 as the backbone and optimized the network using a sigmoid cross-entropy loss. Following the same practice by (Mathis et al., 2018), the original frames were cropped around the mice instead of downsampling.

2D human pose estimation We implemented the Simple-Baseline (Xiao et al., 2018) for its near state-of-the-art performance in 2D human pose estimation with simple architectural designs. This method leverages off-the-shelf object detectors to first locate the candidate subject(s) and per-forms pose estimation over the cropped and resized regions. Compared to DLC/DeeperCut, additional deconvolutional layers are added to the backbone network to generate higher-resolution heatmap outputs.

Multi-view 3D human pose estimation Learnable Triangulation (Iskakov et al., 2019) adopts a similar volumetric approach except that features extracted by a 2D backbone network, instead of raw pixel values, are used to construct the 3D inputs. Similar to SimpleBaseline, a 2D backbone network processes cropped and resized images, where the resulting multi-view features are unprojected on-the-fly to construct the volumetric inputs in the end-to-end training.

Monocular 3D human pose estimation (Pavllo et al., 2019) presented a training scheme for sparsely labeled videos that also leveraged temporal semi-supervision. Instead of using a smoothness constraint, temporal convolutions are performed over sequences of predicted 2D poses obtained from off-the-shelf estimators to regress 3D poses, with additional supervision from a 3D-to-2D backprojection loss and a bone length consistency loss between predictions on labeled and unlabeled frames. Note that we did not specifically train a 3D root joint trajectory model as in the original implementation but directly used the ground truth 3D animal centroids for convenience. Without easy access to off-the-shelf detectors for our keypoint and view set in mice, we employed our best performing 2D model to obtain initial 2D pose estimates.

In addition to the aforementioned approaches, we have adapted a 2D variant of our proposed temporal constraint and applied it to the DLC architecture, similar to Deep-GraphPose (Wu et al., 2020). Instead of using a final sigmoid activation and optimizing against target probability maps, we performed a *soft argmax* on the resulting 2D heatmaps and applied both a supervised regression loss and an unsupervised temporal loss as described in Sects. 3.2 and 3.3, except in the 2D pixel space.

For all approaches, ResNet-50 was used as the backbone network if not otherwise specified. The 2D mouse bounding boxes were computed from 2D projections of ground-truth 3D poses. For 2D approaches, the 2D poses were first estimated separately in each camera view and triangulated into 3D using the same median-based protocol as described in Sect. 3.1. The Protocol 1 MPJPE results were reported for each approach under different annotation conditions (5%, 10%, 50% and 100%).

4.5 Implementation Details

We implemented a standard 3D UNet (Ronneberger et al., 2015) with skip connections to perform our method's 3D pose estimation. The number of feature channels is [64, 64, 128, 128, 256, 256, 512, 512, 256, 256, 128, 128, 64, 64] in the encoder-decoder architecture, followed by a final $1 \times 1 \times 1$ convolution layer outputting one heatmap for each joint position. The encoder consists of four basic blocks with two $3 \times 3 \times 3$ convolution layers with padding 1 and stride 1, one ReLU activation and one $2 \times 2 \times 2$ max pooling for downsampling. The decoder consists of three downsampling blocks, each with one $2 \times 2 \times 2$ transpose convolution layer of stride 2 and two $3 \times 3 \times 3$ convolution layers. The 3D keypoint coordinates were estimated by applying *soft argmax* (Sun et al., 2018) over the predicted heatmaps. We did not explore additional 3D CNN architectures, as this is not the focus of the paper, but we expect that the semi-supervised training strategy should generalize easily to different model

architecture, as demonstrated for 2D in later sections (Section 1).

We trained all models using an Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-7$) with a constant learning rate of 0.0001 for a maximum of 1200 epochs. We used the model checkpoint with the best internal validation MPJPE for evaluation on the test set.

Empirically, we found that a warm-start strategy that only incorporated the unsupervised loss during a later stage performed better for training the temporal+extra models. A similar strategy was also used by Xiong et al. (2021). The temporal+extra models were only supervised by the pose regression loss during the first third of the training epochs, and the unsupervised temporal loss was added afterwards.

5 Results and Discussion

In this section, we quantitatively and qualitatively evaluate the performance gains of our semi-supervised approach.

5.1 Localization Accuracy

We first validated the performance of our semi-supervised approach across 5%, 10%, 50% and 100% annotation conditions using MPJPE and its two variants (Fig.3). Compared to fully supervised models, the temporal consistency constraint generally improved the keypoint localization accuracy, especially in the low annotation conditions. The temporal baseline models improved the MPJPE by 3.0% and 34.8% respectively using 5% and 10% of the training samples. With additional temporal supervision in "temporal+extra" models, our approach improved localization errors by 36.5% and 38.6% for the same low annotation condition.

To confirm that this improvement in localization accuracy could not simply be obtained via post-processing, we tested deliberate smoothing of baseline model predictions using different smoothing methods and window sizes (the full comparisons are presented in "Appendix A"). Despite the obvious decrease in trajectory oscillations from temporal smoothing ("Appendix A" Fig.7), no type of post hoc smoothing improved localization accuracy more than 1%. This suggests that the unsupervised temporal constraint encourages more selective and flexible adaptation of the spatio-temporal features, rather than naive filtering.

5.2 Temporal Smoothness

We first performed a qualitative examination of the movement trajectories of four different keypoint positions over 1000 frames (Fig. 4a). Given the same amount of labeled training data, the temporal approach produced noticeably
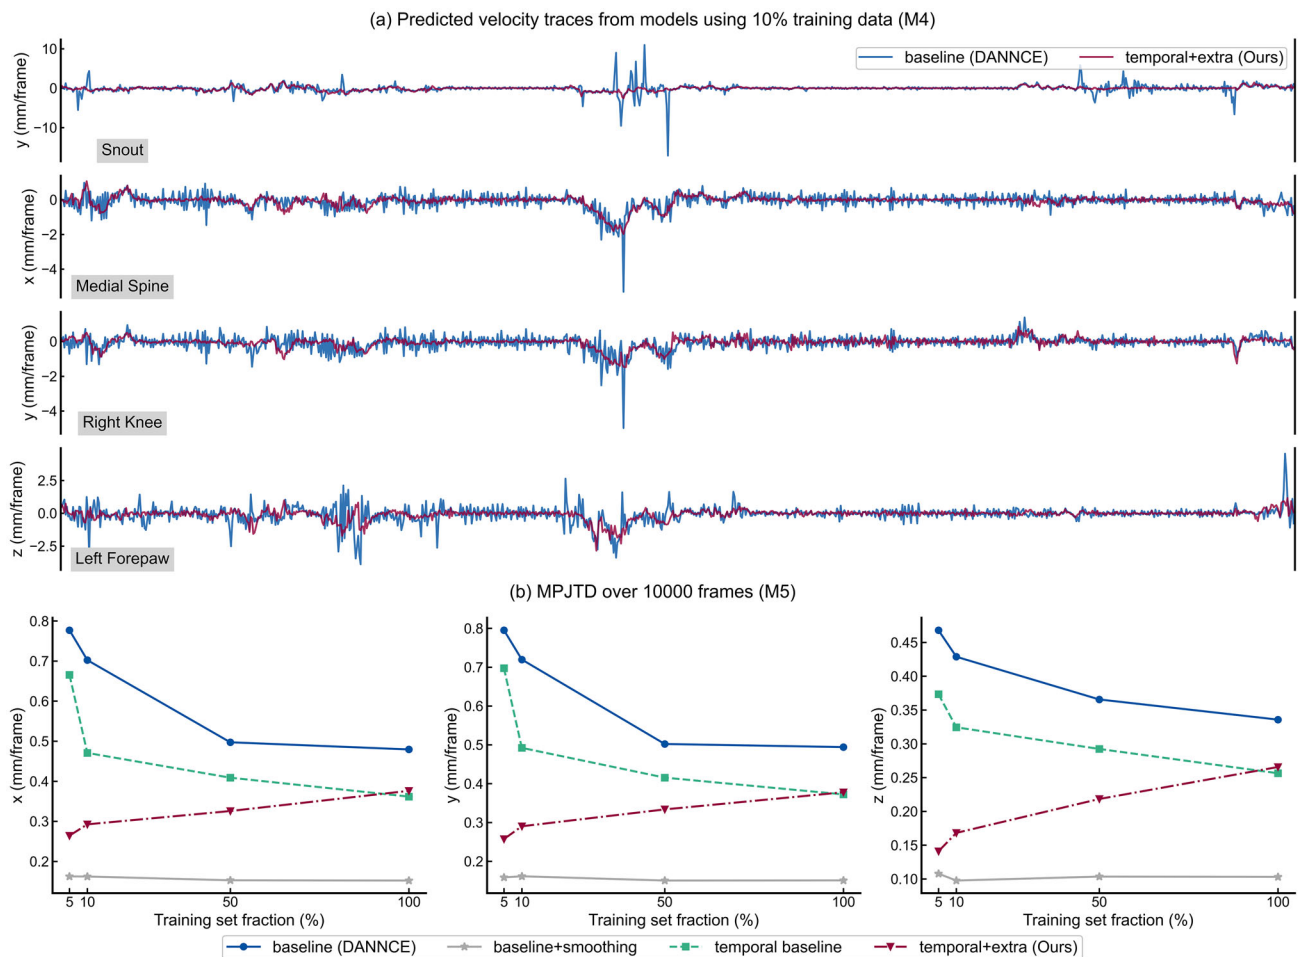
Fig. 4 Analysis of temporal smoothness. **a** Selected coordinate velocities of four different keypoint positions (snout, medial spine, right knee, left forehand) over 1000 consecutive frames from test mouse M4. **b** Quantitative MPJTD results across different training schemes over 10,000 frames from test mouse M5. Our temporal models yield more stable movement trajectories than the baseline fully supervised models

smoother keypoint movement trajectories compared to baseline.

We then quantitatively evaluated MPJTD over a longer period of 10000 frames (Fig. 4b). The inclusion of temporal supervision improved MPJTD by 15.6%, 29.6%, 18.4% and 24.3% for each of the four annotation conditions and by 67.8%, 59.6%, 36.1% and 22.0% when additional unlabeled chunks were added. Post hoc temporal smoothing achieved superior trajectory smoothness as indicated by MPJTD (gray lines), but only resulted in marginal improvement in MPJTD. Meanwhile, the temporal semi-supervised models improved both MPJTD and MPJPE when compared to the baseline models. This reiterates the importance of having a set of comprehensive and complementary performance metrics: MPJTD metric should not be interpreted alone but rather in concert with basic localization accuracy metrics.

5.3 Body Skeleton Consistency

We also quantitatively analyzed the length variations of different body segments of 10,000 consecutive frames (Fig. 5). For simplicity, we grouped the 22 body segments into four general categories: head, trunk, forelimb and hindlimb, and selected two from each category for presentation.

While the fully supervised models struggled to preserve anatomical consistency in low annotation conditions, temporal semi-supervision helped to produce more consistent body structure. The temporal models exhibited less variability in predicted body segment lengths and more closely matched ground-truth average values, especially for the **head** and **trunk**. For body segments with higher coefficients of variation in the ground-truth data (**forelimb, hindlimb**), the addition of temporal supervision generally decreased such variability.
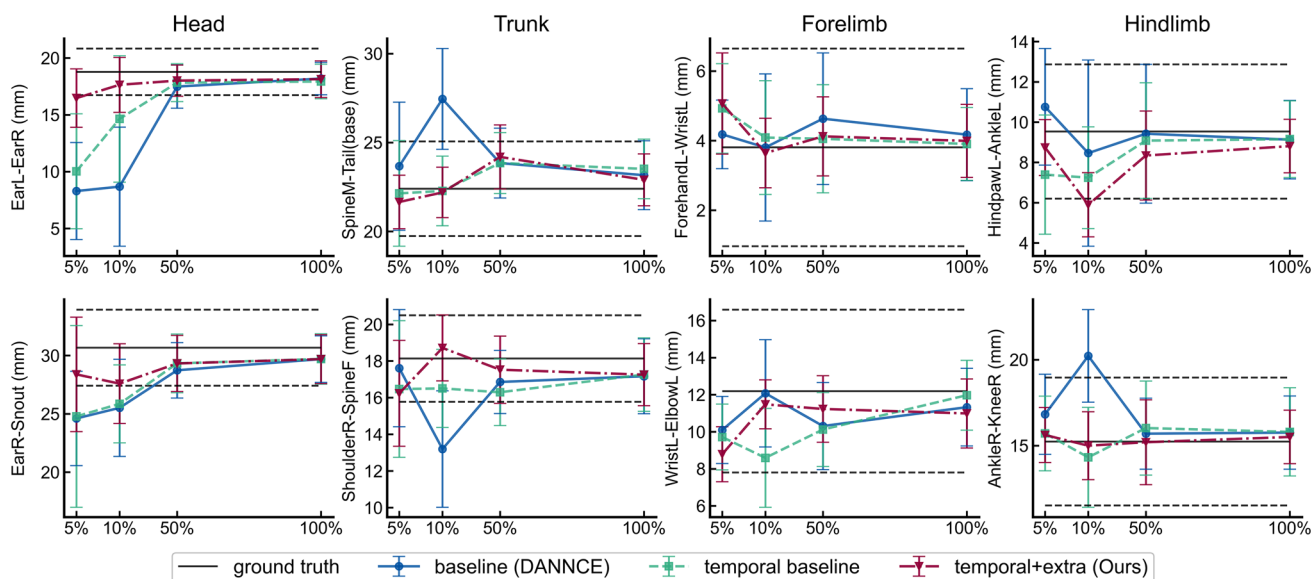
Fig. 5 Body segment length consistency. Plots reporting the statistics of eight different body segment lengths. The solid black horizontal line in each plot represents the mean body segment length computed from man- ually labeled ground-truth, and the horizontal dashed lines encompass corresponding standard deviations. Error bars are standard deviation

5.4 Qualitative Performance on Difficult Poses

In practice, we have identified that baseline models are prone to producing inaccurate keypoint predictions in low annotation regimes, especially for the limbs, when animals are in specific rearing poses. Aside from changes in appearance, such behaviors take place at lower frequencies than others and are thus underrepresented in labeled training data. We therefore also presented qualitative visualization results for one example sequence of rearing behavior frames.

While the baseline 10% model predicted malformed skeletons due to the limited label availability (Fig. 6 blue bounding boxes), the addition of temporal supervision produced marked improvements in physical plausibility. With supervision from additional unlabeled temporal chunks, the "temporal+extra" model produced qualitatively better predictions, even when compared to the 100% baseline model. In cases where the fully supervised baseline model made inaccurate estimates of difficult hindlimb positions (Fig. 6 red bounding box), the semi-supervised approach, with only 10% of the labeled data, better recovered the overall posture.

5.5 Quantitative Comparisons with Other Approaches

We quantitatively examined the proposed method's performance against other widely-adopted animal and human pose estimation approaches, as summarized in Table 1.

Methods for post hoc triangulation of 2D poses Our proposed method consistently outperforms approaches that first independently estimate 2D pose in each camera view and reconstruct the 3D poses via post hoc triangulation. Compared to implicit optimization against heatmap targets, we observed that adapting existing 2D architectures to direct regression of keypoint coordinates effectively improved the overall metric performance (Table 1 "DLC + *soft argmax*"). While approaches like SimpleBaseline appeared sensitive to the quality of 2D bounding boxes, the *soft argmax* approach was able to operate robustly over full-sized images (i.e no cropping or resizing). Applying the 2D variant of our proposed temporal semi-supervision method further improved the performance under all annotation conditions, which implies that the temporal constraint behaves as a powerful prior for recovering plausible poses in both 2D and 3D.

The monocular 3D pose method Considering the inherent ambiguities in monocular 3D representations, it was expected that the precision of monocular estimation would be lower than for multi-view methods. We observed that even with 100% of the training data and temporal regularization, monocular estimation performed worse than many of the multi-view approaches we tested even when using just 5% of the training data. These observations are consistent with what has been reported in previous literature (Iskakov et al., 2019; Bolaños et al., 2021).

Multi-view 3D pose estimation methods We did not observe particular advantages of using 3D volumes constructed from 2D features maps vs. raw pixel values. This
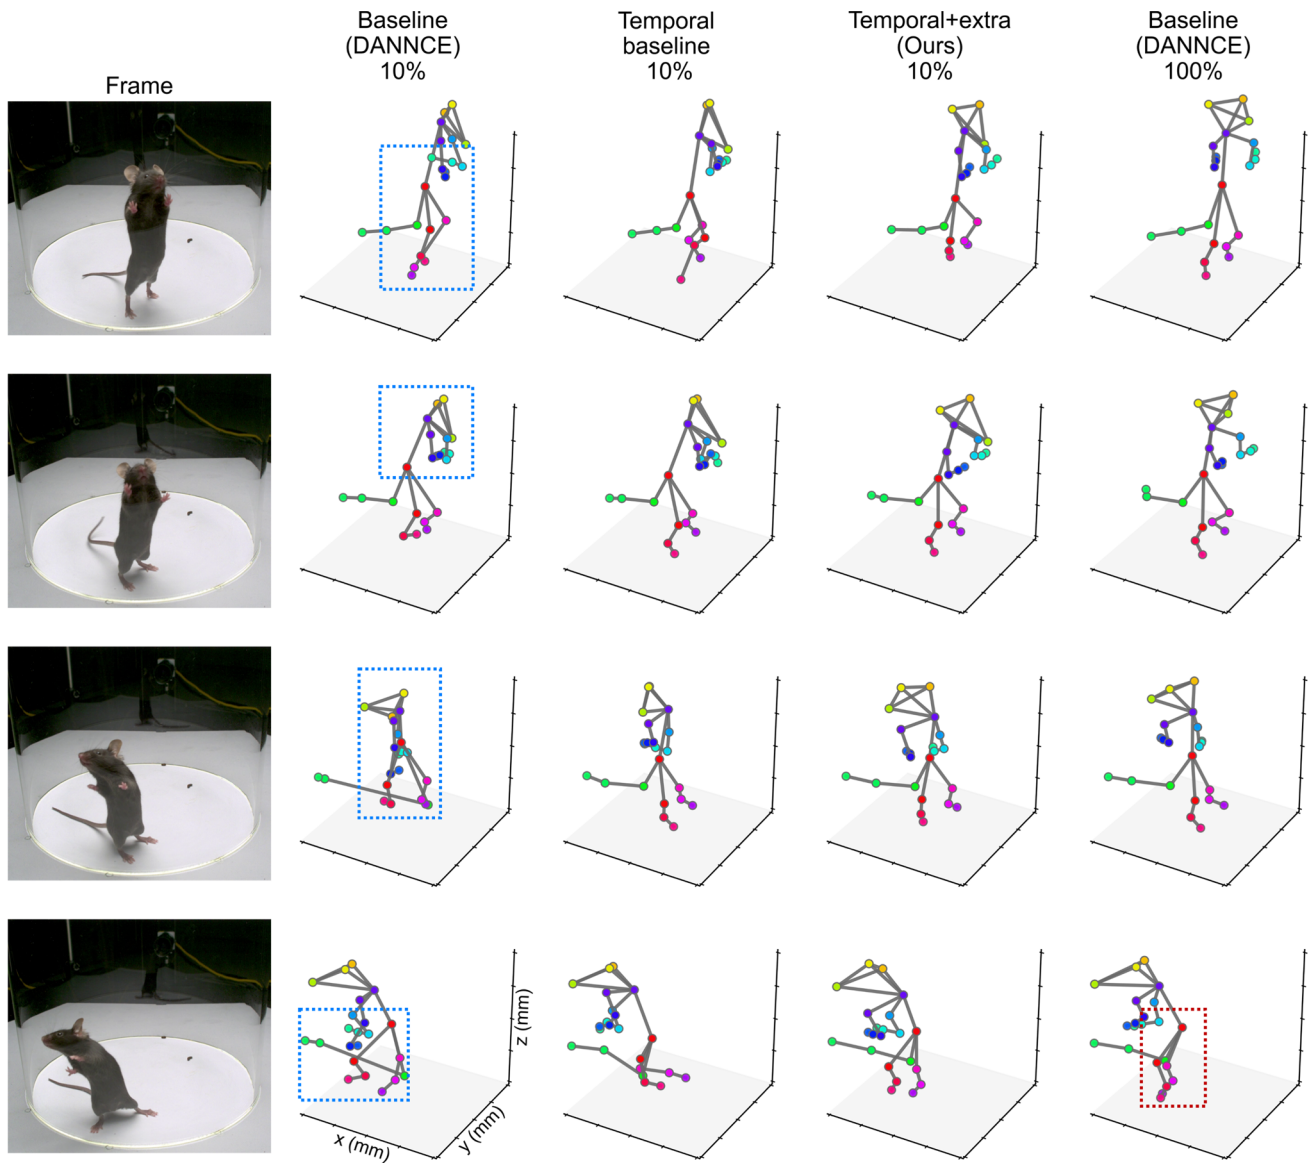
Fig. 6 Qualitative visualization on difficult rearing poses. All 3D visualizations are plotted on the same spatial scale. With 10% of the training samples, the fully supervised baseline model consistently yields inaccurate predictions (blue bounding boxes). Even with 100% of the training samples, the model is still prone to making mistakes on limb landmarks (red bounding box). Many of these errors are corrected via temporal supervision when using just 10% of the labeled data (Color figure online)

likely implies that feature-based volumetric approaches require more accurate 2D feature extraction, via backbone networks pretrained on large-scale 2D pose datasets (Tu et al., 2020). For the human pose case, strong off-the-shelf 2D pose estimators already exist, whereas such options are limited for animal applications. Our results suggest that volume construction directly from pixels, i.e., the strategy used in our temporal semi-supervision method, is the more suitable choice for 3D animal pose estimation in cases where species-specific training data are scarce. This conclusion should

nevertheless be re-evaluated in the future once larger 2D animal pose datasets become available.

6 Conclusion

In this paper, we present a state-of-the-art semi-supervised approach that exploits implicit temporal information to improve the precision and consistency of markerless 3D mouse pose estimation. The approach improves a suite of metrics, each providing a complementary measure of model

performance, and the approach is particularly effective when the labeled data are scarce. Along with the newly released mouse pose dataset, these enhancements will facilitate ongoing efforts to measure freely moving animal behavior across different species and environments.

Supplementary information We provide a video file that demonstrates one example multi-view sequence from the released 3D mouse pose dataset. The original videos are slowed down by 0.3x for better visualization.

Declarations

Appendix A. Additional Quantitative Results

We provide additional metric evaluation results using different smoothing strategies as discussed in section 5.1. We also qualitatively demonstrate the effects of such post hoc smoothing on the original trajectories (see Table 2 and Fig. 7).

Table 2 Complete localization metric comparison

	Baseline	MovAvg 5	MovAvg 10	MovAvg 15	MovAvg 20	MovAvg 25	MovAvg 30	G 5	G 10	G 15	G 20	G 25	G 30	Temp, baseline	Temp. +extra
MPJPE (mm)															
5%	12.87	12.83	12.80	12.79	12.79	12.78	12.78	12.83	12.82	12.81	12.80	12.79	12.78	12.49	8.17
10%	10.90	10.87	10.85	10.81	10.79	10.78	10.78	10.88	10.86	10.83	10.82	10.80	10.79	7.11	6.69
50%	4.99	4.97	4.98	5.00	5.05	5.11	5.19	4.97	4.96	4.96	4.97	5.00	5.03	4.83	5.04
100%	4.36	4.31	4.33	4.35	4.41	4.47	4.56	4.33	4.31	4.31	4.33	4.35	4.38	4.37	4.14
PA-MPJPE (mm)															
5%	10.80	10.77	10.75	10.74	10.74	10.73	10.73	10.78	10.76	10.75	10.74	10.74	10.74	11.07	6.68
10%	9.48	9.43	9.39	9.35	9.33	9.33	9.33	9.44	9.41	9.38	9.36	9.35	9.34	6.49	5.88
50%	4.81	4.78	4.77	4.78	4.80	4.84	4.88	4.79	4.78	4.77	4.77	4.77	4.79	4.65	4.76
100%	4.28	4.23	4.21	4.21	4.23	4.27	4.31	4.25	4.22	4.21	4.20	4.21	4.22	4.28	3.97
N-MPJPE (mm)															
5%	12.70	12.66	12.63	12.62	12.61	12.60	12.59	12.66	12.65	12.63	12.62	12.61	12.61	12.39	8.00
10%	10.84	10.80	10.78	10.74	10.71	10.70	10.69	10.82	10.79	10.77	10.75	10.73	10.72	7.06	6.64
50%	4.99	4.97	4.98	5.01	5.05	5.11	5.18	4.98	4.97	4.97	4.98	5.00	5.03	4.87	5.04
100%	4.39	4.35	4.37	4.38	4.43	4.49	4.57	4.37	4.35	4.35	4.36	4.38	4.40	4.41	4.16

We recorded the changes in MPJPE, PA-MPJPE and N-MPJPE after applying 12 different post hoc smoothing strategies. We used either moving average smoothing ("MovAvg") or Gaussian smoothing ("G"), with window size of 5, 10, 15, 20, 25 or 30 frames
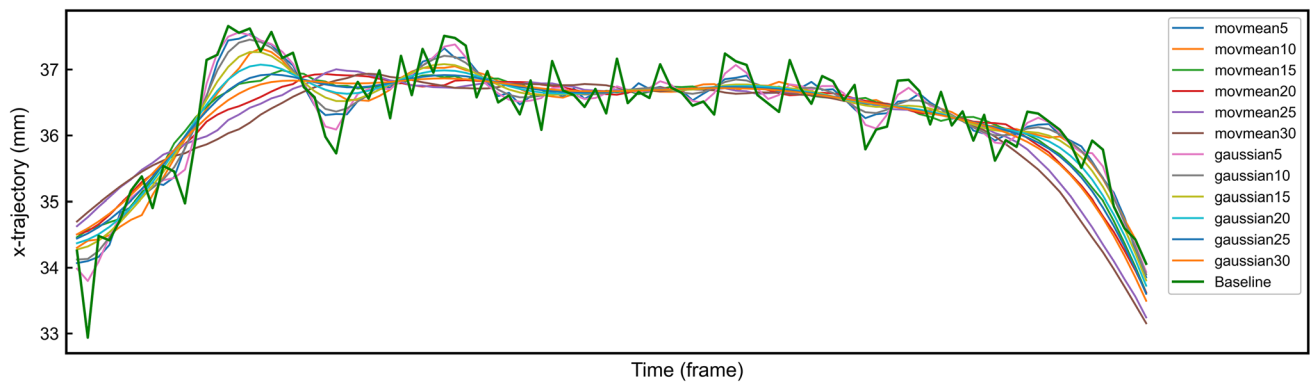
Fig. 7 Visualization of different smoothing strategies. The thick green line corresponds to the original trajectory predicted by the 10% baseline model (Color figure online)

Appendix B. The Multi-view 3D Mouse Pose Dataset

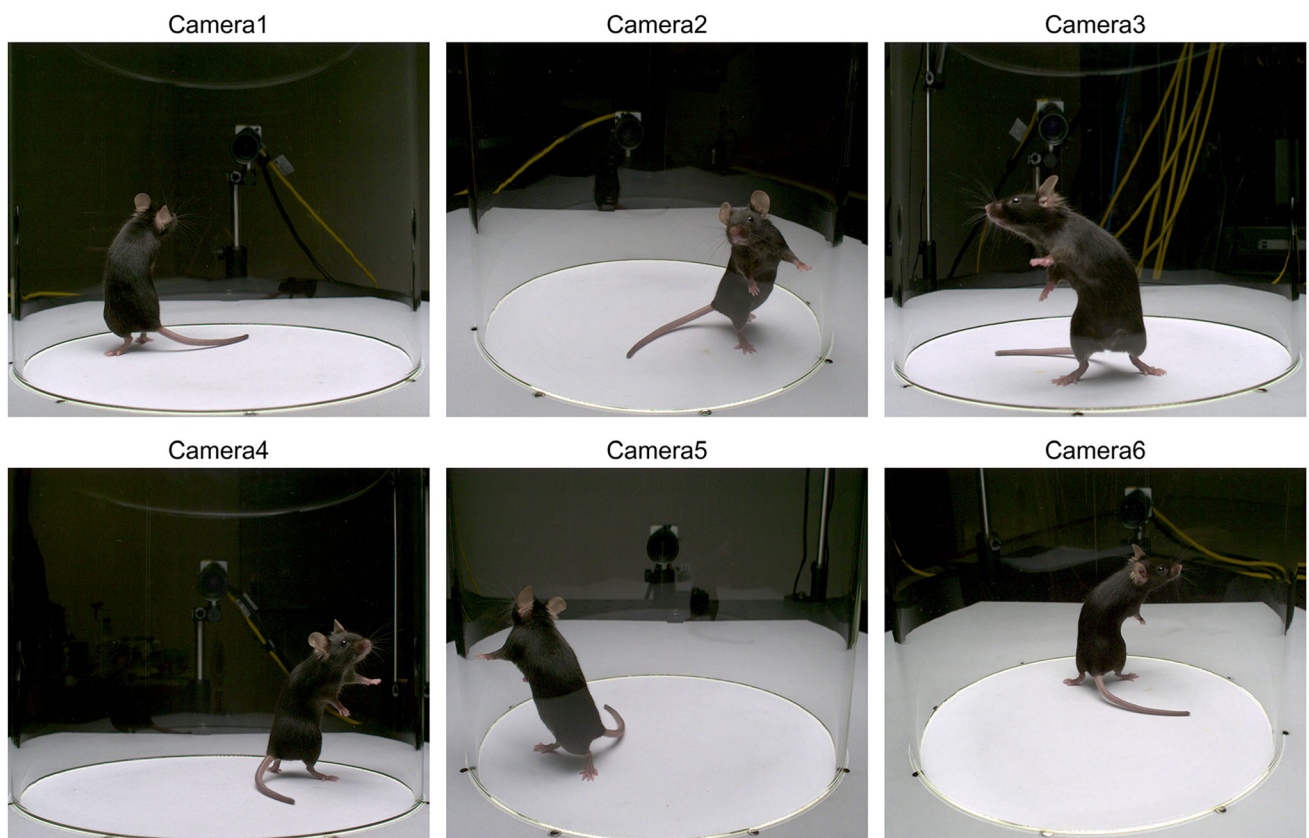We provide supplementary figures that demonstrate the released 3D mouse pose dataset (see Figs. 8 and 9).

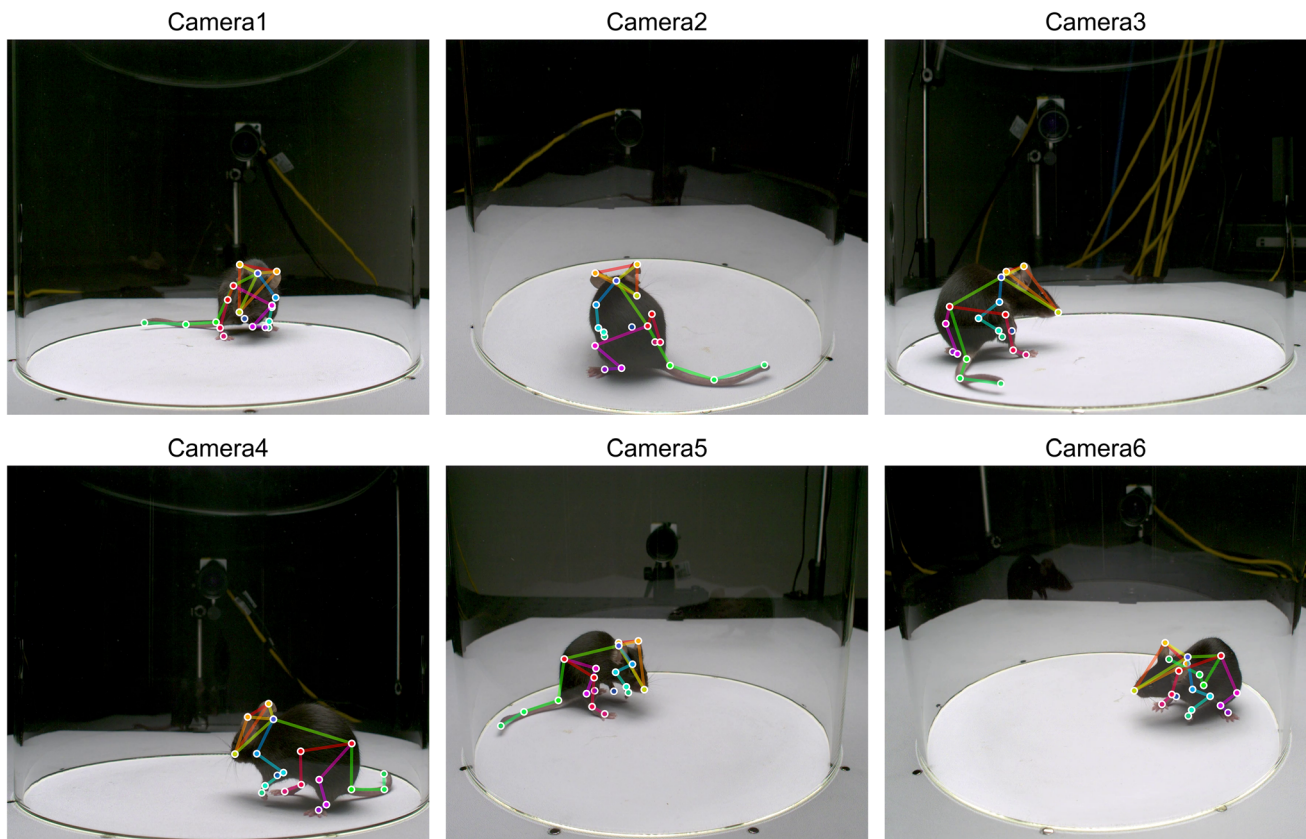**Fig. 8** Multi-view captures from the released mouse dataset

Fig. 9 Multi-view captures from the released mouse dataset (overlaid with ground-truth annotations)

References

Bala, P. C., Eisenreich, B. R., Yoo, S. B. M., Hayden, B. Y., Park, H. S., & Zimmermann, J. (2020). Automated markerless pose estimation in freely moving macaques with OpenMonkeyStudio. *Nature Communications, 11*(1), 1–12.

Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., & Raffel, C. A. (2019). Mixmatch: A holistic approach to semisupervised learning. *Advances in Neural Information Processing Systems, 32*.

Bolaños, L. A., Xiao, D., Ford, N. L., LeDue, J. M., Gupta, P. K., Doebeli, C., & Murphy, T. H. (2021). A three-dimensional virtual mouse generates synthetic training data for behavioral analysis. *Nature Methods, 18*(4), 378–381.

Cao, J., Tang, H., Fang, H.-S., Shen, X., Lu, C., & Tai, Y.-W. (2019). Cross-domain adaptation for animal pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 9498–9507).

Chen, C.-H., Tyagi, A., Agrawal, A., Drover, D., Mv, R., Stojanov, S., Rehg, & J.M. (2019). Unsupervised 3d pose estimation with geometric self-supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5714–5724).

Chen, L., Lin, S.-Y., Xie, Y., Lin, Y.-Y., & Xie, X. (2021). Temporal-aware self-supervised learning for 3d hand pose and mesh estimation in videos. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision* (pp. 1050–1059).

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607).

Dabral, R., Mundhada, A., Kusupati, U., Afaque, S., Sharma, A., & Jain, A. (2018). Learning 3d human pose from structure and motion. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 668–683).

Desmarais, Y., Mottet, D., Slangen, P., & Montesinos, P. (2021). A review of 3d human pose estimation algorithms for markerless motion capture. *Computer Vision and Image Understanding, 212*, 103275.

Dunn, T. W., Marshall, J. D., Severson, K. S., Aldarondo, D. E., Hildebrand, D. G., Chettih, S. N., et al. (2021). Geometric deep learning enables 3d kinematic profiling across species and environments. *Nature Methods, 18*(5), 564–573.

Ellenbroek, B., & Youn, J. (2016). Rodent models in neuroscience research: Is it a rat race? *Disease Models & Mechanisms, 9*(10), 1079–1087.

Gosztolai, A., Günel, S., Lobato-Ríos, V., Pietro Abrate, M., Morales, D., Rhodin, H., & Ramdya, P. (2021). LiftPose3D, a deep learning-based approach for transforming two-dimensional to three-dimensional poses in laboratory animals. *Nature Methods, 18*(8), 975–981.

Günel, S., Rhodin, H., Morales, D., Campagnolo, J., Ramdya, P., & Fua, P. (2019). DeepFly3D, a deep learning-based approach for 3d limb and appendage tracking in tethered, adult Drosophila. *Elife, 8*, e48571.

He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729–9738).

Hossain, M. R. I., & Little, J. J. (2018). Exploiting temporal information for 3d human pose estimation. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 68–84).

Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., & Schiele, B. (2016). Deepercut: A deeper, stronger, and faster multiperson pose estimation model. In *European conference on computer vision* (pp. 34–50).

Ionescu, C., Papava, D., Olaru, V., & Sminchisescu, C. (2013). 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.

Iqbal, U., Molchanov, P., & Kautz, J. (2020). Weakly-supervised 3d human pose learning via multi-view images in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5243–5252).

Iskakov, K., Burkov, E., Lempitsky, V., & Malkov, Y. (2019). Learnable triangulation of human pose. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 7718–7727).

Jain, A., Patel, H., Nagalapatti, L., Gupta, N., Mehta, S., Guttula, S., & Munigala, V. (2020). Overview and importance of data quality for machine learning tasks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 3561–3562).

Joska, D., Clark, L., Muramatsu, N., Jericevich, R., Nicolls, F., Mathis, A., & Patel, A. (2021). AcinoSet: A 3d pose estimation dataset and baseline models for cheetahs in the wild. In *2021 IEEE international conference on robotics and automation (ICRA)* (pp. 13901–13908).

Kar, A., Häne, C., & Malik, J. (2017). Learning a multi-view stereo machine. *Advances in Neural Information Processing Systems, 30*.

Karashchuk, P., Rupp, K. L., Dickinson, E. S., Walling-Bell, S., Sanders, E., Azim, E., & Tuthill, J. C. (2021). Anipose: A toolkit for robust markerless 3d pose estimation. *Cell Reports, 36*(13), 109730.

Kearney, S., Li, W., Parsons, M., Kim, K. I., & Cosker, D. (2020). Rgbd-dog: Predicting canine pose from rgbd sensors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8336–8345).

Kocabas, M., Karagoz, S., & Akbas, E. (2019). Self-supervised learning of 3d human pose using multi-view geometry. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 1077–1086).

Liu, R., Shen, J., Wang, H., Chen, C., Cheung, S.-C., & Asari, V. (2020). Attention mechanism exploits temporal contexts: Real-time 3d human pose reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5064–5073).

Liu, X., Yu, S.-Y., Flierman, N. A., Loyola, S., Kamermans, M., Hoogland, T. M., & De Zeeuw, C. I. (2021). Optiflex: Multiframe animal pose estimation combining deep learning with optical flow. *Frontiers in Cellular Neuroscience, 15*, 621252.

Machado, A. S., Darmohray, D. M., Fayad, J., Marques, H. G., & Carey, M. R. (2015). A quantitative framework for whole-body coordination reveals specific deficits in freely walking ataxic mice. *elife, 4*, e07892.

Marshall, J. D., Klibaite, U., Aldarondo, D. E., Olveczky, B., & Timothy, W. D. (2021). The PAIR-R24M dataset for multi-animal 3d pose estimation. In *Thirty-fifth conference on neural information processing systems datasets and benchmarks track (round 1)*.

Marshall, J. D., Li, T., Wu, J. H., & Dunn, T. W. (2022). Leaving flatland: Advances in 3d behavioral measurement. *Current Opinion in Neurobiology, 73*, 102522.

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). Deeplabcut: Markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience, 21*(9), 1281–1289.

Mimica, B., Dunn, B. A., Tombaz, T., Bojja, V. S., & Whitlock, J. R. (2018). Efficient cortical coding of 3d posture in freely behaving rats. *Science, 362*(6414), 584–589.

Moskvyak, O., Maire, F., Dayoub, F., & Baktashmotlagh, M. (2020). Learning landmark guided embeddings for animal reidentification. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision workshops* (pp. 12–19).

Mu, J., Qiu, W., Hager, G. D., & Yuille, A. L. (2020). Learning from synthetic animals. IN *Proceedings of the IEEE/CVF conference on computervision and pattern recognition* (pp. 12386–12395).

Nibali, A., He, Z., Morgan, S., & Prendergast, L. (2018). Numerical coordinate regression with convolutional neural networks. arXiv preprint arXiv:1801.07372.

Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., & Qu, L. (2017). Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1944–1952).

Pavllo, D., Feichtenhofer, C., Grangier, D., & Auli, M. (2019). 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 7753–7762).

Pereira, T. D., Aldarondo, D. E., Willmore, L., Kislin, M., Wang, S.S.-H., Murthy, M., & Shaevitz, J. W. (2019). Fast animal pose estimation using deep neural networks. *Nature Methods, 16*(1), 117–125.

Pereira, T. D., Tabris, N., Matsliah, A., Turner, D. M., Li, J., Ravindranath, S., Papadoyannis, E. S., Normand, E., Deutsch, D. S., Wang, Z. Y., & McKenzie-Smith, G. C. (2022). Sleap: A deep learning system for multi-animal pose tracking. *Nature Methods, 19*, 486–495.

Reed, G. F., Lynn, F., & Meade, B. D. (2002). Use of coefficient of variation in assessing variability of quantitative assays. *Clinical and Vaccine Immunology, 9*(6), 1235–1239.

Rhodin, H., Salzmann, M., & Fua, P. (2018). Unsupervised geometry-aware representation for 3d human pose estimation. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 750–767).

Rhodin, H., Spörri, J., Katircioglu, I., Constantin, V., Meyer, F., Müller, E., & Fua, P. (2018). Learning monocular 3d human pose estimation from multi-view images. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8437–8446).

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention* (pp. 234–241).

Sarafianos, N., Boteanu, B., Ionescu, B., & Kakadiaris, I. A. (2016). 3d human pose estimation: A review of the literature and analysis of covariates. *Computer Vision and Image Understanding, 152*, 1–20.

Sigal, L., Balan, A. O., & Black, M. J. (2010). Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International Journal of Computer Vision, 87*(1), 4–27.

Spurr, A., Iqbal, U., Molchanov, P., Hilliges, O., & Kautz, J. (2020). Weakly supervised 3d hand pose estimation via biomechanical constraints. In *European conference on computer vision* (pp. 211–228).

Sun, X., Xiao, B., Wei, F., Liang, S., & Wei, Y. (2018). Integral human pose regression. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 529–545).

Tu, H., Wang, C., & Zeng, W. (2020). Voxelpose: Towards multi-camera 3d human pose estimation in wild environment. In *European conference on computer vision* (pp. 197–212).

Wandt, B., & Rosenhahn, B. (2019). Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estima-

tion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 7782–7791).

Wandt, B., Rudolph, M., Zell, P., Rhodin, H., & Rosenhahn, B. (2021). Canonpose: Selfsupervised monocular 3d human pose estimation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 13294–13304).

Wang, J., Yan, S., Xiong, Y., & Lin, D. (2020). Motion guided 3d pose estimation from videos. In *European conference on computer vision* (pp. 764–780).

Wedel, A., Pock, T., Zach, C., Bischof, H., & Cremers, D. (2009). An improved algorithm for TV-L^1 optical flow. In *Statistical and geometrical approaches to visual motion analysis* (pp. 23–45). Springer.

Wu, A., Buchanan, E. K., Whiteway, M., Schartner, M., Meijer, G., Noel, J.-P., et al. (2020). Deep graph pose: A semi-supervised deep graphical model for improved animal pose tracking. *Advances in Neural Information Processing Systems, 33*, 6040–6052.

Xiao, B., Wu, H., & Wei, Y. (2018). Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 466–481).

Xiong, B., Fan, H., Grauman, K., & Feichtenhofer, C. (2021). Multiview pseudo-labeling for semi-supervised learning from video. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 7209–7219).

Yao, Y., Jafarian, Y., & Park, H.S. (2019). Monet: Multiview semi-supervised keypoint detection via epipolar divergence. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 753–762).

Zhang, L., Dunn, T., Marshall, J., Olveczky, B., & Linderman, S. (2021). Animal pose estimation from video data with a hierarchical von Mises–Fisher–Gaussian model. In *International conference on artificial intelligence and statistics* (pp. 2800–2808).

Zimmermann, C., Schneider, A., Alyahyay, M., Brox, T., & Diester, I. (2020). Freipose: a deep learning framework for precise animal motion capture in 3d spaces. BioRxiv.