

X-MoGen: Unified Motion Generation Across Humans and Animals

Xuan Wang^{1*}, Kai Ruan^{3*}, Liyang Qian¹, Guo Zhi Zhi², Chang Su¹, Gaoang Wang^{1†}

¹Zhejiang University

²Institute of Artificial Intelligence (TeleAI), China Telecom

³Renmin University of China

xuanw@zju.edu.cn, kairuan@ruc.edu.cn, gaoangwang@intl.zju.edu.cn

Abstract

Text-driven motion generation has attracted increasing attention due to its broad applications in virtual reality, animation, and robotics. While existing methods typically model human and animal motion separately, a joint cross-species approach offers key advantages, such as a unified representation and improved generalization. However, morphological differences across species remain a key challenge, often compromising motion plausibility. To address this, we propose **X-MoGen**, the first unified framework for cross-species text-driven motion generation covering both humans and animals. X-MoGen adopts a two-stage architecture. First, a conditional graph variational autoencoder learns canonical T-pose priors, while an autoencoder encodes motion into a shared latent space regularized by morphological loss. In the second stage, we perform masked motion modeling to generate motion embeddings conditioned on textual descriptions. During training, a morphological consistency module is employed to promote skeletal plausibility across species. To support unified modeling, we construct **UniMo4D**, a large-scale dataset of 115 species and 119k motion sequences, which integrates human and animal motions under a shared skeletal topology for joint training. Extensive experiments on UniMo4D demonstrate that X-MoGen outperforms state-of-the-art methods on both seen and unseen species.

1 Introduction

Motion generation aims to synthesize high-quality, physically plausible 3D motion sequences. It has become an active research area with wide-ranging applications in film production, video games, virtual/augmented reality, and robotics (Zhu et al. 2023; Sahili, Neji, and Tabia 2025). Various control modalities have been explored for motion generation, including text (Guo et al. 2022), music (Siyao et al. 2022), sketch-based input (Zhong et al. 2025), and multi-modal signals (Guo et al. 2025; Zhang et al. 2024b). Of these modalities, text has emerged as a particularly popular interface due to its intuitive and expressive nature. Despite recent advances, most text-driven motion generation methods focus exclusively on either human (Guo et al. 2022; Hong

*These authors contributed equally.

†Corresponding author.

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




Figure 1: X-MoGen achieves a wide range of capabilities within a single unified framework, including generating both human and animal motions from text descriptions and enabling smooth cross-species motion transitions.

et al. 2025) or animal (Wang et al. 2025) motion, treating them as independent tasks. However, unified modeling offers two key advantages (Yang et al. 2024a): (1) Unified representation learning, where aligning motion semantics across species leads to more structured and expressive latent spaces; and (2) Cross-species generalization, as diverse body morphologies and motion patterns improve transferability to unseen species and novel actions.

Previous works typically struggle to achieve unified modeling due to incompatible skeletal structures across datasets, necessitating separate models for each dataset (Guo et al. 2024; Hong et al. 2025). To address this, we introduce **UniMo4D**, a large-scale dataset covering humans and 114 animal species. UniMo4D standardizes skeletal topology while retaining species-specific morphology, facilitating

learning in a shared representation space. This provides a foundation for joint human-animal motion modeling. However, unifying such diverse species introduces new challenges, as variations in morphology and motion dynamics make it difficult to generate consistent and realistic cross-species motions with a single model.

To address these challenges, we propose **X-MoGen**, a unified framework designed for human-animal motion generation. As illustrated in Figure 1, X-MoGen synthesizes motions for both humans and animals from a single text description, providing flexible and generalizable control. It also supports applications such as cross-species transformation by producing smooth and semantically consistent transitions, such as transforming a tiger into a human. X-MoGen operates in two stages. The first stage uses a conditional graph variational autoencoder to generate species-specific canonical T-pose priors, while a motion autoencoder encodes motion sequences into a compact latent space regularized by a morphological loss. In the second stage, motion is synthesized using a masked modeling mechanism (Li et al. 2023, 2024b) conditioned on the input text. To maintain structural coherence, a morphological consistency module predicts skeletal lengths from the latent representation and incorporates them into the loss computation.

We evaluate X-MoGen on the UniMo4D dataset against several state-of-the-art baselines. Results show that our model achieves strong performance on both seen and unseen species, demonstrating its ability to generalize across diverse body morphologies. Quantitative results further confirm that X-MoGen generates realistic motions that are consistent with species-specific morphology and dynamics.

Our main contributions are summarized as follows:

- To the best of our knowledge, **X-MoGen** is the first model designed for text-driven motion generation across humans and animals, addressing challenges in morphological variation and cross-species generalization.
- We construct **UniMo4D**, a large-scale benchmark that combines human and animal motion data with a unified skeletal topology, enabling joint training across species.
- Extensive experiments demonstrate that our method achieves strong performance on both seen and unseen species, supported by both qualitative and quantitative evaluations.

2 Related Work

2.1 Text-Driven Human Motion Generation

Text-driven motion generation methods can be broadly categorized into quantization-based and non-quantization-based approaches. Quantization-based methods (Jiang et al. 2023; Zhang et al. 2023a; Guo et al. 2024; Zhong et al. 2023; Pinyoanuntapong et al. 2024; Liao et al. 2025) discretize motion sequences using various quantization schemes and predict motion tokens from text. However, this discretization often leads to information loss. Alternatively, non-quantization-based approaches operate directly in the continuous motion space (Rombach et al. 2022; Tevet et al. 2023; Zhang et al. 2023b, 2024a; Meng et al. 2025b; Hong

et al. 2025). For example, MDM (Tevet et al. 2023) employs a diffusion model on raw motion sequences for synthesis. Subsequent works leverage latent diffusion models (Rombach et al. 2022) to improve generation quality while reducing computational cost. MARDM (Meng et al. 2025b) introduces a masked autoregressive framework for high-fidelity motion generation. However, such models are typically trained on skeletons with fixed bone lengths. The recent ShapeMove (Liao et al. 2025) supports variable bone lengths by learning additional shape features, allowing skeletal morphology to be controlled by text. In contrast, X-MoGen learns species-specific motions without relying on additional shape features, directly adapting to skeletal variations present in the data.

2.2 Text-Driven Animal Motion Generation

Generating animal motions from text is inherently more complex than human motion generation, due to the vast diversity in animal morphology and behavior. OmniMotionGPT (Yang et al. 2024b) models both human and animal motions, enabling animal synthesis via transfer learning from human motion-text pairs. Crucially, human motion serves as a key condition, guiding animal motion generation within the same category. Motion Avatar (Zhang et al. 2024c) employs large language models as high-level planners to generate 3D animal avatar meshes. AniMo (Wang et al. 2025) introduces spatiotemporal modeling and species-aware modulation for text-driven animal motion generation. However, these methods do not explicitly address morphological variations across species. In contrast, our approach incorporates T-pose encoding and morphology-aware losses to ensure that synthesized motions accurately reflect species-specific structures.

2.3 Unified Human-Animal Motion Modeling

Building a unified framework for both human and animal motion remains challenging. WalkTheDog (Li et al. 2024a) learns a shared phase manifold for unsupervised semantic and temporal alignment across morphologically different subjects, enabling applications such as motion transfer without requiring explicit skeleton correspondences. Any-Top (Gat et al. 2025) generates natural motions aligned with a character’s biomechanical characteristics by conditioning on its skeletal structure. However, the generated motions are sampled from the training distribution and remain relatively uncontrollable. UniMoGen (Khani et al. 2025) synthesizes motions based on the input character’s style, trajectory, and motion history. In contrast to methods requiring user-specified skeletons, our X-MoGen enables cross-species motion generation solely from free-form text.

3 Method

We address the challenge of cross-species motion generation using X-MoGen, a novel two-stage framework illustrated in Figure 2. In the first stage, a Conditional Graph Variational Autoencoder (CGAE) generates canonical T-pose priors conditioned on species, while an Autoencoder (AE) encodes motion sequences into a compact latent space. In the




Figure 2: **Overview of the X-MoGen architecture.** Our two-stage framework first learns a T-pose prior and a compact latent motion space via a Conditional Graph Autoencoder (CGAE) and an Autoencoder (AE). In the second stage, a Masked Transformer (M-Trans) conditions a diffusion model to generate motion from noise, guided by the text description and species-specific T-pose priors produced by the CGAE. An auxiliary MCM promotes morphological constraints in the generated motions.

second stage, a Masked Transformer (M-Trans) synthesizes motion from text, conditioned on T-pose priors derived from species cues in the input. The output is further refined by a diffusion-based generation head. To ensure anatomical plausibility across species, a Morphological Consistency Module (MCM) enforces structural coherence.

3.1 Stage 1: Feature Modeling

T-pose Modeling via CGAE. We propose a Conditional Graph Variational Autoencoder (CGAE) to model species-specific skeletal morphology and generalize to unseen species. CGAE integrates a conditional variational framework (Sohn, Lee, and Yan 2015) with graph convolutional encoders and decoders (Kipf and Welling 2016), explicitly treating the skeletal structure as a graph. CGAE is trained on bone-length vectors $\mathbf{b} \in \mathbb{R}^{N_b}$ extracted from T-poses using *Bone-Length Extraction*. It learns a latent distribution conditioned on species-level semantic embeddings \mathbf{c} derived from BioCLIP-2’s text encoder (Gu et al. 2025). Both encoder and decoder operate on a fixed-topology graph, with bone lengths as node features.

The training objective balances reconstruction accuracy and latent regularization using the evidence lower bound (ELBO) (Kingma, Welling et al. 2013):

$$\mathcal{L}_{\text{CGAE}} = \left\| \hat{\mathbf{b}} - \mathbf{b} \right\|_2^2 + \beta D_{\text{KL}}(q_\phi(\mathbf{z}_{\text{pose}} | \mathbf{b}, \mathbf{c}) \| p(\mathbf{z}_{\text{pose}})), \quad (1)$$

where $\hat{\mathbf{b}}$ is the reconstructed bone-length vector, $p(\mathbf{z}_{\text{pose}}) =$

$\mathcal{N}(\mathbf{0}, \mathbf{I})$ is a standard Gaussian prior, $D_{\text{KL}}(\cdot \| \cdot)$ denotes the Kullback–Leibler divergence that regularizes the posterior, and β is a weighting factor controlling the trade-off between reconstruction fidelity and latent smoothness.

Once trained, CGAE can generate canonical T-poses for arbitrary quadruped species. Given a condition vector \mathbf{c} and a latent sample $\mathbf{z}_{\text{pose}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the decoder predicts a bone-length vector $\hat{\mathbf{b}}$. Final joint positions \mathbf{P} are recovered via *Forward Kinematics* based on a predefined kinematic tree (Parent 2012).

Motion Representation via AE. To encode motion sequences into an efficient and structured representation (Rombach et al. 2022), we employ an Autoencoder (AE) (Bank, Koenigstein, and Giryes 2023) that operates on raw motion inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_L]$. The encoder maps \mathbf{X} to a temporally downsampled latent sequence $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_T] \in \mathbb{R}^{T \times d}$, where $T = \lfloor L/4 \rfloor$ and d denotes the latent dimension. The encoder consists of convolutional layers with temporal strides and residual connections (He et al. 2016), while the decoder uses transposed convolutions to reconstruct the original sequence.

We train the AE by minimizing a weighted combination of reconstruction and morphological loss:

$$\mathcal{L}_{\text{AE}} = \lambda_{\text{MSE}}^{\text{recon}} + \lambda_{\text{morph}}^{\text{recon}} \cdot \mathcal{L}_{\text{morph}}^{\text{recon}}. \quad (2)$$

The reconstruction loss is defined as:

$$\mathcal{L}_{\text{MSE}}^{\text{recon}} = \frac{1}{L} \sum_{t=1}^L \|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2, \quad (3)$$

where $\hat{\mathbf{x}}_t$ is the reconstructed motion at time t .

The morphological loss enforces structural fidelity by comparing bone lengths computed from motions:

$$\mathcal{L}_{\text{morph}}^{\text{recon}} = \frac{1}{L} \sum_{t=1}^L \|\mathcal{B}(\hat{\mathbf{x}}_t) - \mathcal{B}(\mathbf{x}_t)\|_2^2, \quad (4)$$

where $\mathcal{B}(\cdot)$ extracts bone lengths via pairwise joint distances.

3.2 Stage 2: Masked Motion Generation

Recent studies have demonstrated the effectiveness of masked modeling (Li et al. 2023; Guo et al. 2024; Li et al. 2024b; Meng et al. 2025b). Building on this, we adopt a Masked Transformer (M-Trans) with a diffusion head to synthesize motion sequences. The model learns to reconstruct masked motion tokens through a denoising process, guided by text description.

Masked Motion Modeling. M-Trans serves as a context-aware backbone that encodes semantically rich representations for motion generation. Given a latent motion sequence $\mathbf{Z} \in \mathbb{R}^{T \times d}$, a random subset of temporal positions $\mathcal{M} \subset \{1, \dots, T\}$ is masked by replacing the corresponding tokens with a learnable masked token [M]. The canonical T-pose \mathbf{P} is projected to an embedding $\mathbf{p} \in \mathbb{R}^d$ via linear transformation and concatenated with the masked sequence as the first token, forming the input sequence. The query tensor is computed as:

$$\mathbf{Q} = \text{Proj}_Q([\mathbf{p}; \mathbf{Z}_{\text{masked}}]). \quad (5)$$

Textual guidance is extracted from CLIP’s text encoder (Radford et al. 2021), which provides a sentence-level feature \mathbf{s} and word-level features \mathbf{W} . These features are projected to obtain keys and values for cross-attention:

$$\mathbf{K} = \text{Proj}_K([\mathbf{s}; \mathbf{W}]), \quad \mathbf{V} = \text{Proj}_V([\mathbf{s}; \mathbf{W}]). \quad (6)$$

The Transformer fuses motion and text features (Vaswani et al. 2017), producing a contextualized representation $\mathbf{H} \in \mathbb{R}^{(T+1) \times d}$:

$$\mathbf{H} = \text{Transformer}(\mathbf{Q}, \mathbf{K}, \mathbf{V}). \quad (7)$$

Diffusion-based Completion Head. The contextualized representation \mathbf{H} from the M-Trans backbone serves as the conditioning signal for a diffusion-based completion head. This head is realized by a set of Multi-Layer Perceptrons (MLPs), which function as the velocity network \mathbf{v}_θ . For each masked position $i \in \mathcal{M}$, its corresponding contextualized embedding \mathbf{h}_i is extracted from \mathbf{H} . The velocity network is optimized to predict the deterministic velocity field that maps noise $\hat{\mathbf{z}}_i$ to the ground-truth data \mathbf{z}_i , following the flow-matching paradigm introduced by SiT (Ma et al. 2024). The optimization uses the following loss, conditioned on \mathbf{h}_i :

$$\mathcal{L}_{\text{flow}} = \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathbb{E}_{\tau, \hat{\mathbf{z}}_i, \mathbf{z}_i} [\|\mathbf{v}_\theta(\mathbf{z}_{\tau, i}, \tau, \mathbf{h}_i) - (\mathbf{z}_i - \hat{\mathbf{z}}_i)\|_2^2], \quad (8)$$

where $\mathbf{z}_{\tau, i} = (1 - \tau)\hat{\mathbf{z}}_i + \tau\mathbf{z}_i$ is the interpolated latent state at time τ .

Morphological Consistency Module. To encourage realistic and consistent skeletal structures in generated motions, we design the Morphological Consistency Module (MCM), denoted as f_{MCM} . This module serves as a morphology-aware regularizer that maps the latent motion sequence $\hat{\mathbf{Z}}$ to a static vector representing the skeleton’s bone lengths. Specifically, a Gated Recurrent Unit (GRU) (Chung et al. 2014) encodes the temporal dynamics of $\hat{\mathbf{Z}}$, and its final hidden state is processed by a MLP to predict the bone-length vector. The MCM is pre-trained on the latent space of AE and remains frozen during motion generation. To ensure morphological consistency, we introduce a guidance loss that penalizes deviations from a canonical bone-length vector \mathbf{b} , computed from the ground-truth bone lengths:

$$\mathcal{L}_{\text{morph}}^{\text{guide}} = \left\| f_{\text{MCM}}(\hat{\mathbf{Z}}) - \mathbf{b} \right\|_2^2. \quad (9)$$

Gradients from this loss are backpropagated only through the masked positions \mathcal{M} in $\hat{\mathbf{Z}}$, while unmasked positions are detached.

Training Objective. The total training loss combines the flow objective and morphological consistency:

$$\mathcal{L}_{\text{gen}} = \mathcal{L}_{\text{flow}} + \lambda_{\text{morph}}^{\text{guide}} \cdot \mathcal{L}_{\text{morph}}^{\text{guide}}. \quad (10)$$

3.3 Inference

During inference, the latent sequence $\mathbf{Z}^{(0)} \in \mathbb{R}^{T \times d}$ is initialized with a learnable masked token [M] at all positions, where the superscript (0) denotes the initial iteration step. The full sequence is generated over R iterative steps. At each iteration step r , a subset of the masked positions is selected for refinement. For these positions, latent embeddings are generated via a diffusion-based denoising process. Specifically, Gaussian noise $\hat{\mathbf{z}}_i$ is initialized for each selected position i , and N denoising steps are performed using the diffusion head, conditioned on the text prompt and the unmasked tokens generated so far.

To improve text alignment, classifier-free guidance (Ho and Salimans 2022) is applied at each denoising step to compute the guided velocity:

$$\mathbf{v}_{\text{guided}, i}^{(n)} = \mathbf{v}_{\text{uncond}, i}^{(n)} + \omega \cdot \left(\mathbf{v}_{\text{cond}, i}^{(n)} - \mathbf{v}_{\text{uncond}, i}^{(n)} \right), \quad (11)$$

where $\omega > 0$ is the guidance scale, and the superscript (n) denotes the n -th denoising step within each iteration.

After denoising, the generated latent embeddings \mathbf{z}_i replace the corresponding [M] positions in $\mathbf{Z}^{(r)}$, yielding the updated sequence $\mathbf{Z}^{(r+1)}$. This process repeats until all positions are filled. Finally, the completed latent sequence \mathbf{Z} is decoded by the AE to produce the motion output \mathbf{X} .

4 Experiments

Dataset. To support joint modeling of human and animal motions, we construct UniMo4D, a unified dataset that integrates HumanML3D (Guo et al. 2022), KIT-ML (Plappert,

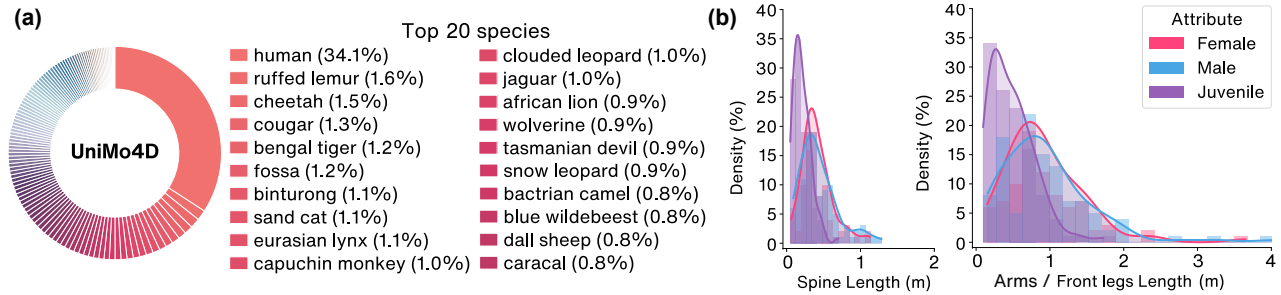


Figure 3: **Statistics of the UniMo4D dataset.** (a) Species distribution. (b) Length distribution of key bones.

Methods	Venue	R-Precision \uparrow			FID \downarrow	MM-Dist \downarrow	Diversity \rightarrow	MME \downarrow
		Top-1 \uparrow	Top-2 \uparrow	Top-3 \uparrow				
Real	–	0.899 \pm .001	0.972 \pm .001	0.989 \pm .001	0.000 \pm .000	0.470 \pm .133	19.407 \pm .255	-
MDM	ICLR'23	0.283 \pm .007	0.444 \pm .010	0.556 \pm .010	4.684 \pm .833	7.766 \pm .109	18.199 \pm .250	0.313 \pm .000
T2M-GPT	CVPR'23	0.382 \pm .002	0.463 \pm .002	0.677 \pm .002	5.058 \pm .059	5.196 \pm .014	18.452 \pm .316	0.409 \pm .000
AtT2M	ICCV'23	0.334 \pm .001	0.466 \pm .001	0.695 \pm .001	5.895 \pm .192	5.695 \pm .030	15.895 \pm .219	0.456 \pm .000
MMM	CVPR'24	0.585 \pm .001	0.746 \pm .001	0.822 \pm .001	5.200 \pm .027	3.742 \pm .004	18.225 \pm .284	0.390 \pm .000
MoMask	CVPR'24	0.554 \pm .002	0.752 \pm .002	0.845 \pm .001	0.745 \pm .018	3.602 \pm .006	19.151 \pm .285	0.220 \pm .000
MARDM	CVPR'25	0.823 \pm .001	0.927 \pm .001	0.960 \pm .001	0.189 \pm .002	0.874 \pm .001	18.706 \pm .423	0.251 \pm .000
AniMo	CVPR'25	0.596 \pm .003	0.786 \pm .002	0.869 \pm .002	0.533 \pm .007	3.174 \pm .009	17.955 \pm .305	0.211 \pm .000
X-MoGen	Ours	0.848 \pm .001	0.941 \pm .001	0.968 \pm .001	0.050 \pm .001	0.742 \pm .002	19.332 \pm .480	0.201 \pm .000

Table 1: Comparison of text-driven motion generation on the UniMo4D dataset. The right arrow (\rightarrow) indicates that values closer to the ground truth motion are better. The best results for each metric are shown in bold.

Mandery, and Asfour 2016), and AniMo4D (Wang et al. 2025). Among various motion representations (Guo et al. 2022; Lu et al. 2023; Meng et al. 2025b,a; Xiao et al. 2025), we adopt the *essential feature groups* (Meng et al. 2025b), a compact and effective scheme. Each frame is represented as a vector of dimension $N_j \times 3 + 1$, where N_j denotes the number of joints. The root joint features include rotation velocity, linear velocity, and height (y-axis), while the remaining joints are represented by 3D positions. To create a unified representation, we first standardize the skeletal structures from all datasets. Specifically, we remap joints from KIT-ML and reorder those from AniMo4D to align with the HumanML3D convention (Petrovich, Black, and Varol 2023; Chen and UniMocap 2023), scaling bone lengths accordingly. All skeletons are then standardized to a common 25-joint topology, which incorporates three virtual tail joints to accommodate diverse morphologies. For tailless species (e.g., humans), these virtual joints are initialized to be collocated with the pelvic root. This neutral configuration ensures a consistent data structure and prevents motion artifacts, without compromising the fidelity of core body movements (Yang et al. 2023). As a result, every motion sequence in UniMo4D is represented by a skeleton with $N_j = 25$ joints and a consistent topology.

Motion sequences with lengths $L \in (18, 300)$ are retained, resulting in 118,663 sequences spanning 115 species. Figure 3 (a) presents the species distribution, and Figure 3 (b) illustrates bone length variations across species, high-

lighting the challenge of unified modeling across diverse morphologies.

In our experiments, five species comprising 5,179 motion sequences are held out as an unseen-species test set. The remaining species are split into training, validation, and test sets with proportions of 0.80, 0.05, and 0.15, respectively.

Evaluation Metrics. A text-motion matching model is re-trained on the UniMo4D dataset to serve as the evaluation backbone (Guo et al. 2022; Meng et al. 2025b). Based on its learned feature space, motion quality, semantic alignment, diversity, and morphological consistency are evaluated. **Frchet Inception Distance (FID)** assesses motion realism by comparing the feature distributions of generated and real motions. **R-Precision** measures text-motion alignment through retrieval accuracy, while **Multimodal Distance (MM-Dist)** quantifies the average distance between motion and text embeddings. **Diversity** measures variability across motions generated from different prompts.

To evaluate morphological plausibility, **Mean Morphological Error (MME)** is introduced to measure skeletal consistency over time. Specifically, it quantifies deviations in bone lengths from a reference skeleton, where $\hat{\mathbf{b}}_t \in \mathbb{R}^{N_b}$ denotes the bone lengths at frame t , and $\mathbf{b} \in \mathbb{R}^{N_b}$ is the corresponding reference vector. MME is defined as the average L1 distance over the sequence:

$$\text{MME} = \frac{1}{L \cdot N_b} \sum_{t=1}^L \|\hat{\mathbf{b}}_t - \mathbf{b}\|_1 \quad (12)$$

Methods	Venue	R-Precision \uparrow			FID \downarrow	MM-Dist \downarrow	Diversity \rightarrow	MME \downarrow
		Top-1 \uparrow	Top-2 \uparrow	Top-3 \uparrow				
Real	—	0.594 \pm .003	0.826 \pm .003	0.924 \pm .002	0.000 \pm .000	0.310 \pm .001	20.015 \pm .399	-
MDM	ICLR'23	0.099 \pm .003	0.186 \pm .003	0.268 \pm .004	28.481 \pm 1.678	11.490 \pm .076	23.587 \pm .652	0.310 \pm .002
T2M-GPT	CVPR'23	0.101 \pm .001	0.184 \pm .001	0.256 \pm .001	55.775 \pm .223	11.025 \pm .022	13.729 \pm .267	0.457 \pm .000
AttT2M	ICCV'23	0.093 \pm .001	0.180 \pm .001	0.254 \pm .002	65.045 \pm 1.172	15.279 \pm .051	16.125 \pm .252	0.485 \pm .000
MMM	CVPR'24	0.113 \pm .003	0.213 \pm .001	0.298 \pm .002	49.840 \pm .242	10.489 \pm .025	13.887 \pm .227	0.445 \pm .000
MoMask	CVPR'24	0.138 \pm .002	0.261 \pm .003	0.369 \pm .003	24.441 \pm .193	8.180 \pm .020	18.383 \pm .248	0.235 \pm .000
MARDM	CVPR'25	0.143 \pm .002	0.265 \pm .001	0.374 \pm .002	34.124 \pm .250	7.690 \pm .004	13.779 \pm .305	0.301 \pm .000
AniMo	CVPR'25	0.114 \pm .002	0.213 \pm .002	0.304 \pm .002	45.605 \pm .245	9.688 \pm .016	15.594 \pm .353	0.238 \pm .000
X-MoGen	Ours	0.148 \pm .002	0.271 \pm .002	0.372 \pm .002	19.935 \pm .093	7.939 \pm .010	17.652 \pm .230	0.229 \pm .000

Table 2: Comparison of text-driven motion generation on the UniMo4D unseen species test dataset. The right arrow (\rightarrow) indicates that values closer to the ground truth motion are better. The best results for each metric are shown in bold.

where L is the sequence length and N_b is the number of bones. Lower MME reflects more consistent species-specific morphology, reducing unrealistic bone deformations.

Baselines. To the best of our knowledge, X-MoGen is the first unified framework designed for cross-species motion generation involving both humans and animals. As no existing method is specifically designed for this setting, we compare our approach against recent state-of-the-art models developed for either human or animal motion generation. Specifically, we evaluate several human motion baselines, including MDM (Tevet et al. 2023), T2M-GPT (Zhang et al. 2023a), AttT2M (Zhong et al. 2023), MMM (Pinyoanuntapong et al. 2024), MoMask (Guo et al. 2024), and MARDM (Meng et al. 2025b), along with AniMo (Wang et al. 2025), which focuses on animal motion. All baseline models are retrained from scratch on the UniMo4D dataset using their official implementations and original data formats (Xiao et al. 2025; Meng et al. 2025a). Although these methods use different internal data formats, we extract essential motion features from their outputs to enable a fair and consistent evaluation (Meng et al. 2025b).

4.1 Quantitative Results

We evaluate X-MoGen on the UniMo4D dataset. All experiments are repeated 10 times, and we report mean values with 95% confidence intervals (Guo et al. 2024; Wang et al. 2025). Table 1 shows the results on the test set. X-MoGen outperforms all baselines with significantly lower FID and higher R-Precision, highlighting the superior motion quality and text alignment of our method. Table 2 reports performance on the unseen species test set, which includes species excluded during training and presents a more challenging generalization scenario. X-MoGen maintains strong performance across Top-1, Top-2, FID, and MME, which are the primary metrics for evaluating motion plausibility and morphological consistency. Table 3 compares several representative motion compression methods for reconstruction. Among quantization-based approaches, VQ (Van Den Oord, Vinyals et al. 2017), FSQ (Mentzer et al. 2023), and RVQ (Lee et al. 2022; Zeghidour et al. 2021) all struggle

to capture the complexity of the UniMo4D dataset, although FSQ and RVQ offer moderate improvements over vanilla VQ. Our AE does not use quantization and achieves much better reconstruction accuracy. This shows that continuous latent modeling works better for diverse motion data.

4.2 Qualitative Results

To further assess the effectiveness of our model, we qualitatively compare X-MoGen with three strong baselines: MoMask, MARDM, and AniMo, selected based on their quantitative performance. Figure 4 presents two representative examples. In the human motion case, X-MoGen accurately captures the prompt containing “running” and “cartwheel”, generating a coherent sequence with a smooth transition into the cartwheel. In contrast, MoMask reverses the intended action order, while MARDM and AniMo produce only the cartwheel, omitting the preceding context. Notably, the virtual tail nodes introduced for topological consistency do not degrade motion realism across all methods. In the unseen species case, X-MoGen generates motion that is both morphologically consistent and semantically faithful to the prompt. MoMask fails by abruptly morphing the animal into a human, MARDM misinterprets the prompt and produces unstable body rotations, and AniMo incorrectly generates a left-turn motion. These results demonstrate that X-MoGen effectively unifies human and animal motion synthesis, exhibiting strong generalization across species.

4.3 Ablation Study

We conduct ablation experiments to evaluate the contribution of each component. The ablated elements include: (1) word-level features \mathbf{W} from CLIP; (2) the canonical T-pose prior from CGAE; (3) the morphological reconstruction loss $\mathcal{L}_{\text{morph}}^{\text{recon}}$ in the AE; and (4) the morphological guidance loss $\mathcal{L}_{\text{morph}}^{\text{guide}}$ in the MCM.

As shown in Table 4, removing word-level embeddings slightly increases FID, indicating that fine-grained text features enhance semantic alignment. Excluding the CGAE prior raises both MME and FID, highlighting its role in producing structurally consistent T-pose priors. While remov-

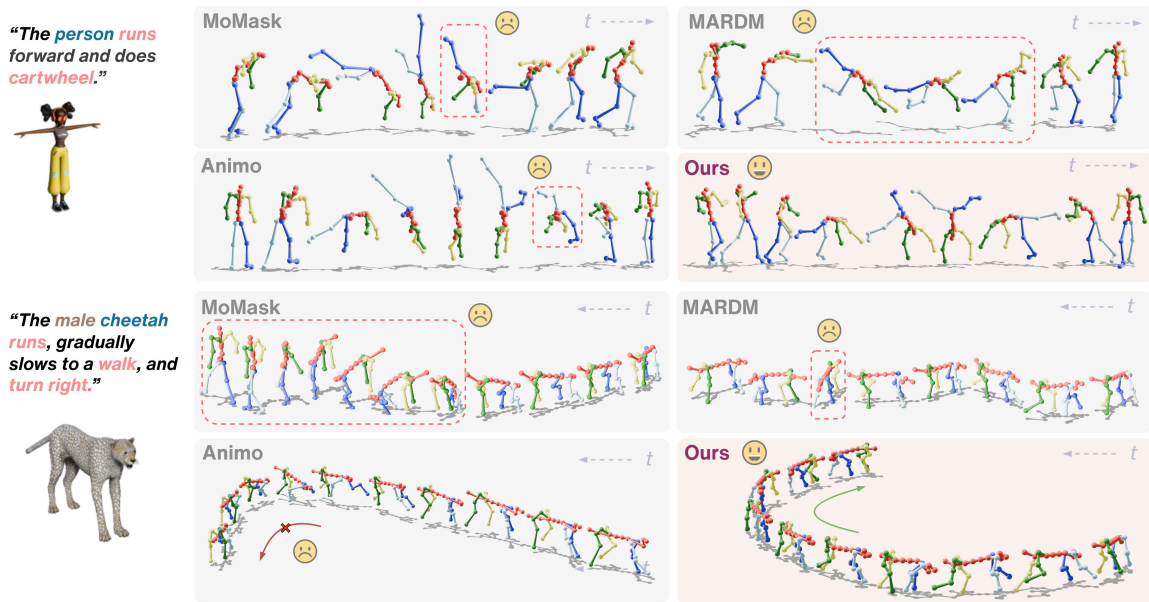


Figure 4: Qualitative results for human and unseen animal motion. Red dashed boxes and red arrows highlight implausible motion artifacts.

ing MCM leads to a marginal increase in FID (+0.001), it notably worsens MME, suggesting its importance for motion consistency. Similarly, ablating the AE reconstruction loss also results in a performance drop.

These results validate the necessity of each component and demonstrate their complementary roles in achieving high-quality, semantically aligned motion generation.

Compressor	Top-1 \uparrow	FID \downarrow	MM-Dist \downarrow	MME \downarrow
VQ (T2M-GPT)	0.581	5.010	4.015	0.389
FSQ (MMM)	0.629	4.272	3.414	0.388
RVQ (MoMask)	0.700	0.451	2.299	0.220
AE (X-MoGen)	0.898	0.008	0.496	0.064

Table 3: Quantitative results of different motion compressors on the UniMo4D test set.

Methods	Top-1 \uparrow	FID \downarrow	MM-Dist \downarrow	MME \downarrow
X-MoGen	0.848	0.050	0.742	0.201
w/o W	0.835	0.053	0.806	0.205
w/o CGAE	0.841	0.055	0.758	0.228
w/o $\mathcal{L}_{\text{morph}}^{\text{recon}}$	0.843	0.108	0.845	0.237
w/o MCM	0.848	0.049	0.743	0.238

Table 4: Ablation study on the UniMo4D test set.

4.4 Applications

Universal Motion Generation. X-MoGen presents a unified framework for generating both human and animal motions from text within a shared representation space. The

framework supports humans and 114 animal species, leveraging CGAE-generated, species-specific T-poses to generalize to new quadrupeds without additional training. This enables intuitive AI-driven motion creation and supports diverse applications including content production, biology-inspired simulation, digital storytelling, and interactive media (Abootorabi et al. 2025).

Cross-Species Motion Transformation. X-MoGen enables smooth motion transitions between species. Given two motion sequences (e.g., a tiger walking and a human walking), the AE encoder extracts latent features, and mask tokens [M] are inserted at the transition point. M-Trans, with the diffusion head and AE decoder, generates realistic cross-species transitions. As shown in Figure 1, this enables smooth morphing between animal and human motions, supporting seamless cross-species motion editing and completion, valuable for animation and automated shapeshifting in film (Akber et al. 2023; Tessler et al. 2024).

5 Conclusion

We introduced **X-MoGen**, a unified framework for generating motions across both human and animal domains. X-MoGen models skeletal structures with a CGAE and compresses motions into a latent space with an AE. It further incorporates masked motion modeling conditioned on text and a T-pose prior derived from a CGAE. During training, we leverage MCM to encourage morphologically plausible generation. Additionally, we constructed the **UniMo4D** dataset, which standardizes human and animal motion representations, facilitating future research in cross-species motion generation. Extensive experiments demonstrate that X-MoGen outperforms strong baselines, generating realistic motions with generalization to unseen species.

Acknowledgments

This work was supported by the Zhejiang Provincial Natural Science Foundation of China (No. LZ24F030005), the National Natural Science Foundation of China (No. 62576308), the Fundamental Research Funds for the Central Universities (No. 226-2025-00167), and the Scientific Research Foundation of Sichuan Provincial Department of Science and Technology, China (No. 2024YFHZ0001).

References

- Abootorabi, M. M.; Ghahroodi, O.; Zahraei, P. S.; Behzadasl, H.; Mirrokni, A.; Salimipannah, M.; Rasouli, A.; Behzadipour, B.; Azarnoush, S.; Maleki, B.; et al. 2025. Generative AI for Character Animation: A Comprehensive Survey of Techniques, Applications, and Future Directions. *arXiv preprint arXiv:2504.19056*.
- Akber, S. M. A.; Kazmi, S. N.; Mohsin, S. M.; and Szczesna, A. 2023. Deep learning-based motion style transfer tools, techniques and future challenges. *Sensors*, 23(5): 2597.
- Bank, D.; Koenigstein, N.; and Giryas, R. 2023. Autoencoders. *Machine learning for data science handbook: data mining and knowledge discovery handbook*, 353–374.
- Chen, L.-H.; and UniMocap, C. 2023. UniMocap: Unifier for BABEL, HumanML3D, and KIT. <https://github.com/LinghaoChan/UniMoCap>.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Gat, I.; Raab, S.; Tevet, G.; Reshef, Y.; Bermano, A. H.; and Cohen-Or, D. 2025. Anytop: Character animation diffusion with any topology. *arXiv preprint arXiv:2502.17327*.
- Gu, J.; Stevens, S.; Campolongo, E. G.; Thompson, M. J.; Zhang, N.; Wu, J.; Kopanov, A.; Mai, Z.; White, A. E.; Balhoff, J.; et al. 2025. Bioclip 2: Emergent properties from scaling hierarchical contrastive learning. *arXiv preprint arXiv:2505.23883*.
- Guo, C.; Mu, Y.; Javed, M. G.; Wang, S.; and Cheng, L. 2024. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1900–1910.
- Guo, C.; Zou, S.; Zuo, X.; Wang, S.; Ji, W.; Li, X.; and Cheng, L. 2022. Generating Diverse and Natural 3D Human Motions From Text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5152–5161.
- Guo, Z.; Hu, Z.; Zhao, N.; and Soh, D. W. 2025. MotionLab: Unified Human Motion Generation and Editing via the Motion-Condition-Motion Paradigm. *arXiv preprint arXiv:2502.02358*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Hong, S.; Kim, C.; Yoon, S.; Nam, J.; Cha, S.; and Noh, J. 2025. SALAD: Skeleton-aware Latent Diffusion for Text-driven Motion Generation and Editing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 7158–7168.
- Jiang, B.; Chen, X.; Liu, W.; Yu, J.; Yu, G.; and Chen, T. 2023. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36: 20067–20079.
- Khani, A.; Rampini, A.; Atherton, E.; and Roy, B. 2025. UniMoGen: Universal Motion Generation. *arXiv preprint arXiv:2505.21837*.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11523–11532.
- Li, P.; Starke, S.; Ye, Y.; and Sorkine-Hornung, O. 2024a. WalkTheDog: Cross-Morphology Motion Alignment via Phase Manifolds. In *ACM SIGGRAPH 2024 Conference Papers*, 1–10.
- Li, T.; Chang, H.; Mishra, S.; Zhang, H.; Katabi, D.; and Krishnan, D. 2023. Mage: Masked generative encoder to unify representation learning and image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2142–2152.
- Li, T.; Tian, Y.; Li, H.; Deng, M.; and He, K. 2024b. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems*, 37: 56424–56445.
- Liao, T.-H.; Zhou, Y.; Shen, Y.; Huang, C.-H. P.; Mitra, S.; Huang, J.-B.; and Bhattacharya, U. 2025. Shape my moves: Text-driven shape-aware synthesis of human motions. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 1917–1928.
- Lu, S.; Chen, L.-H.; Zeng, A.; Lin, J.; Zhang, R.; Zhang, L.; and Shum, H.-Y. 2023. Humantomato: Text-aligned whole-body motion generation. *arXiv preprint arXiv:2310.12978*.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, 23–40. Springer.
- Meng, Z.; Han, Z.; Peng, X.; Xie, Y.; and Jiang, H. 2025a. Absolute Coordinates Make Motion Generation Easy. *arXiv preprint arXiv:2505.19377*.
- Meng, Z.; Xie, Y.; Peng, X.; Han, Z.; and Jiang, H. 2025b. Rethinking Diffusion for Text-Driven Human Motion Generation: Redundant Representations, Evaluation, and Masked Autoregression. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 27859–27871.

- Mentzer, F.; Minnen, D.; Agustsson, E.; and Tschannen, M. 2023. Finite scalar quantization: Vq-vae made simple. *arXiv preprint arXiv:2309.15505*.
- Parent, R. 2012. *Computer animation: algorithms and techniques*. Newnes.
- Petrovich, M.; Black, M. J.; and Varol, G. 2023. TMR: Text-to-Motion Retrieval Using Contrastive 3D Human Motion Synthesis. *arXiv preprint*.
- Pinyoanuntapong, E.; Wang, P.; Lee, M.; and Chen, C. 2024. Mmm: Generative masked motion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1546–1555.
- Plappert, M.; Mandery, C.; and Asfour, T. 2016. The KIT Motion-Language Dataset. *Big Data*, 4(4): 236–252.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sahili, A. R.; Neji, N.; and Tabia, H. 2025. Text-driven Motion Generation: Overview, Challenges and Directions. *arXiv preprint arXiv:2505.09379*.
- Siyao, L.; Yu, W.; Gu, T.; Lin, C.; Wang, Q.; Qian, C.; Loy, C. C.; and Liu, Z. 2022. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11050–11059.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28.
- Tessler, C.; Guo, Y.; Nabati, O.; Chechik, G.; and Peng, X. B. 2024. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 43(6): 1–21.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-or, D.; and Bermano, A. H. 2023. Human Motion Diffusion Model. In *The Eleventh International Conference on Learning Representations*.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, X.; Ruan, K.; Zhang, X.; and Wang, G. 2025. AniMo: Species-Aware Model for Text-Driven Animal Motion Generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 1929–1939.
- Xiao, L.; Lu, S.; Pi, H.; Fan, K.; Pan, L.; Zhou, Y.; Feng, Z.; Zhou, X.; Peng, S.; and Wang, J. 2025. MotionStreamer: Streaming Motion Generation via Diffusion-based Autoregressive Model in Causal Latent Space. *arXiv preprint arXiv:2503.15451*.
- Yang, J.; Zeng, A.; Zhang, R.; and Zhang, L. 2024a. X-pose: Detecting any keypoints. In *European Conference on Computer Vision*, 249–268. Springer.
- Yang, S.; Wang, Z.; Wu, Z.; Li, M.; Zhang, Z.; Huang, Q.; Hao, L.; Xu, S.; Wu, X.; Yang, C.; et al. 2023. Unifiedgesture: A unified gesture synthesis model for multiple skeletons. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1033–1044.
- Yang, Z.; Zhou, M.; Shan, M.; Wen, B.; Xuan, Z.; Hill, M.; Bai, J.; Qi, G.-J.; and Wang, Y. 2024b. OmnimotionGPT: animal motion generation with limited data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1249–1259.
- Zeghidour, N.; Luebs, A.; Omran, A.; Skoglund, J.; and Tagliasacchi, M. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 495–507.
- Zhang, J.; Zhang, Y.; Cun, X.; Zhang, Y.; Zhao, H.; Lu, H.; Shen, X.; and Shan, Y. 2023a. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14730–14740.
- Zhang, M.; Cai, Z.; Pan, L.; Hong, F.; Guo, X.; Yang, L.; and Liu, Z. 2024a. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6): 4115–4128.
- Zhang, M.; Jin, D.; Gu, C.; Hong, F.; Cai, Z.; Huang, J.; Zhang, C.; Guo, X.; Yang, L.; He, Y.; et al. 2024b. Large motion model for unified multi-modal motion generation. In *European Conference on Computer Vision*, 397–421. Springer.
- Zhang, M.; Li, H.; Cai, Z.; Ren, J.; Yang, L.; and Liu, Z. 2023b. Finemogen: Fine-grained spatio-temporal motion generation and editing. *Advances in Neural Information Processing Systems*, 36: 13981–13992.
- Zhang, Z.; Wang, Y.; Wu, B.; Chen, S.; Zhang, Z.; Huang, S.; Zhang, W.; Fang, M.; Chen, L.; and Zhao, Y. 2024c. Motion avatar: Generate human and animal avatars with arbitrary motion. *arXiv preprint arXiv:2405.11286*.
- Zhong, C.; Hu, L.; Zhang, Z.; and Xia, S. 2023. Att2m: Text-driven human motion generation with multi-perspective attention mechanism. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 509–519.
- Zhong, L.; Guo, C.; Xie, Y.; Wang, J.; and Li, C. 2025. Sketch2Anim: Towards Transferring Sketch Storyboards into 3D Animation. *arXiv preprint arXiv:2504.19189*.
- Zhu, W.; Ma, X.; Ro, D.; Ci, H.; Zhang, J.; Shi, J.; Gao, F.; Tian, Q.; and Wang, Y. 2023. Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 2430–2449.