

A Unified Masked Autoencoder with Patchified Skeletons for Motion Synthesis

Esteve Valls Mascaró¹, Hyemin Ahn², Dongheui Lee^{1,3},

¹ Technische Universität Wien (TUWien)

² Ulsan National Institute of Science & Technology (UNIST)

³ German Aerospace Center (DLR)

{esteve.valls.mascaro, dongheui.lee}@tuwien.ac.at, hyemin.ahn@unist.ac.kr

Abstract

The synthesis of human motion has traditionally been addressed through task-dependent models that focus on specific challenges, such as predicting future motions or filling in intermediate poses conditioned on known key-poses. In this paper, we present a novel task-independent model called UNIMASK-M, which can effectively address these challenges using a unified architecture. Our model obtains comparable or better performance than the state-of-the-art in each field. Inspired by Vision Transformers (ViTs), our UNIMASK-M model decomposes a human pose into body parts to leverage the spatio-temporal relationships existing in human motion. Moreover, we reformulate various pose-conditioned motion synthesis tasks as a reconstruction problem with different masking patterns given as input. By explicitly informing our model about the masked joints, our UNIMASK-M becomes more robust to occlusions. Experimental results show that our model successfully forecasts human motion on the Human3.6M dataset while achieving state-of-the-art results in motion inbetweening on the LaFAN1 dataset for long transition periods.

Introduction

Synthesizing plausible human motions is a long-standing challenge in the computer vision and graphics community. Due to the high complexity of the field, researchers often focus on addressing only a specific subset of human motion synthesis, as the different tasks involve varied input-output patterns for the model. Unlike previous works, we propose a unified architecture for solving various motion synthesis tasks. These different tasks are depicted in Fig. 1 and represent the motivation for the research addressed by this paper.

For instance, the 3D human motion forecasting task aims to predict the human’s subsequent 3D body trajectory based only on past observations. In this field, researchers restrict their model to learning in a causality manner (Jain et al. 2016; Aksan et al. 2021; Valls Mascaró et al. 2022), where the new pose only depends on the previous frames. However, these causal-based approaches are not well-suited for the motion inbetweening task, which involves completing the body poses between the past observations and a known future pose. Most inbetweening models (Duan et al. 2021;

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

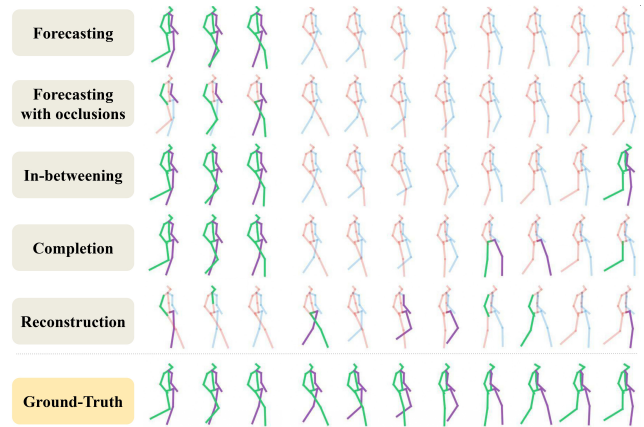


Figure 1: Unified architecture for different human motion synthesis tasks. Green and purple skeletons denote a known skeleton joint, while light red and green represents our model prediction over a masked joint.

Kim et al. 2022; Oreshkin et al. 2022) usually leverage bidirectional temporal relationships since both past and future poses are key for filling the middle poses. However, these works consider the full-body skeleton as a unique condition, constraining the generation process. In our work, we additionally consider the motion completion task, where only independent body parts constrain the motion synthesis in the future, which enhances the freedom of the generation. Finally, motion reconstruction is the task of recovering a given motion subjected to occlusions, as might occur in real-world applications (Bütepage, Kjellström, and Kragic 2018). Dissecting the motion synthesis field in all these different tasks diminishes the versatility of the models. To overcome this challenge, generative approaches (Cai et al. 2021; Tevet et al. 2022) have been proposed recently to tackle motion synthesis with a unified stochastic model, but require high computational resources and are unsuitable for real-time prediction (Tevet et al. 2022) or underperform compared to deterministic models (Cai et al. 2021), as also indicated by Ahn, Valls Mascaró, and Lee (2023).

In this work, we consider both efficiency and versatility as essential characteristics of a model, which should achieve high performance across multiple tasks. Inspired by

the recent success of masked autoencoders in reconstructing images (He et al. 2021; Bao, Dong, and Wei 2021), and videos (Tong et al. 2022), we adapt the motion synthesis as a reconstruction problem: to recover a sequence of masked human skeletons regardless of the masking pattern. In addition, we propose a deconstruction model to decompose human skeletons into body-part patches (i.e., legs, arms, trunk), which boosts the model’s performance while giving more flexibility. Thanks to our deconstruction approach, our model can synthesize motion by leveraging the spatio-temporal relationships of partial body parts through self-attention (Vaswani et al. 2017). Therefore, our **UNified Masked Autoencoder with patchified SKeletons for Motion synthesis**, called **UNIMASK-M**, serves as the first model to address the deterministic motion synthesis task as a reconstruction problem through pose decomposition and masked autoencoders and achieves comparable performance to the state-of-the-art in several motion synthesis tasks. The contributions of our paper can be summarized in three-fold:

1. A unique model that achieves performance comparable to the state-of-the-art in various motion synthesis tasks.
2. A pose decomposition approach that deconstructs a single skeleton into patches, thus giving more freedom in the pose conditioning while boosting the performance.
3. An efficient model suitable for real-time human motion synthesis and robust to occlusions.

Related Work

Literature has addressed each human motion synthesis task as a distinct challenge. This section first provides an overview of each motion synthesis task. Then, we describe the Masked Autoencoders which are the core of our model.

Human Motion Synthesis

Motion synthesis conditioned by key-poses serves as a basis for understanding observed humans’ intentions or enabling agents or robots to collaborate with them.

Motion Forecasting. The aim is to predict future human motion based only on past observations. Early works (Fragkiadaki et al. 2015; Jain et al. 2016) employed Recurrent Neural Networks (RNNs) to process the time-series dependencies of poses. Subsequently, (Mao et al. 2019; Ma et al. 2022) proposed Discrete Cosine Transformations (DCT) to encode temporal information and adopt graph structures to model the spatial relationships. Recent works (Valls Mascaro et al. 2022; Mao, Liu, and Salzmann 2020) adopt the Transformer-based model (Vaswani et al. 2017) to leverage the self-attention among joint information in both space and time. Recently, Guo et al. (2023) used temporal Multi-Layer Perceptrons (MLP) in an autoregressive process, and Xu et al. (2023) proposed to maintain motion equivariance for better motion predictions. However, none of these models can deal with any motion synthesis task if future key-poses condition the motion.

Motion Inbetweening. In this task, the model is given a sequence of past poses and a future target pose and the goal is to fill in the motion between these key-poses. Early studies (Holden, Saito, and Komura 2016) used RNNs to model

the temporal relationships between the skeletons. Harvey et al. (2020b) introduced scheduled noise injection in the target pose to encourage motion variations. Duan et al. (2021) adopted a self-attention block to refine the interpolated poses and a key-pose embedding to support multiple whole pose conditions in a sequence. More recently, Oreshkin et al. (2022) adapted cross-attention to synthesize the missing frames. Duan et al. (2021); Oreshkin et al. (2022) referenced the predicted motion to an interpolated sequence between key-poses, thus reformulating the inbetweening problem as the synthesis of the variation of the poses from an interpolated reference. Kim et al. (2022) highlights the motivation of conditioning the motion synthesis through partial body poses but still requires the last pose to be full-body for sequence completion. Finally, Qin, Zheng, and Zhou (2022) shows the benefit of refining the synthesized motion with self-attention, and applies post-processing steps and complex losses to boost the performance. Inspired by the image inpainting strategy using Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020), Tevet et al. (2022) evaluates their model for human motion completion conditioned by key-poses. However, their work requires high computation, and the denoising process impedes real-time execution. Unlike other works, we propose an efficient model deals with motion inbetweening and completion, independently of having the totality of the key-poses.

Motion Reconstruction. In real scenarios, the estimated human poses can be subjected to occlusions. However, the aforementioned models overlook this problem by projecting the full-body skeletons into a single representation space. Therefore, the model cannot effectively handle partial occlusions within a single pose, just the complete absence of the full skeleton. Jiang, Chen, and Guo (2022); Zhu et al. (2022); Baradel et al. (2022) adopt masked modeling for motion reconstruction in 3D human pose estimation. However, Jiang, Chen, and Guo (2022); Zhu et al. (2022) cannot handle partial pose masking, and Baradel et al. (2022) uses joint-level masking, which we show that underperforms.

Unified Motion Synthesis. Embracing motion synthesis only accounting for particular masking patterns allows researchers to design models that can excel at one of the specific tasks. However, designing a unified architecture that performs well across multiple tasks of motion synthesis is challenging. To tackle this challenge, Cai et al. (2021) uses a Conditional Variational Autoencoder (CVAE) (Kingma, Welling et al. 2019) to sample from the masked sequence distribution and condition the generated missing poses by the anchor poses. Tevet et al. (2022) evaluated their DDPM architecture in different motion synthesis tasks. However, we prioritize deterministic over stochastic approaches as we observe that diversity in motions can be obtained by adding noise to the anchor poses (Harvey et al. 2020b), while having closer predictions to reality is essential for real scenarios.

Semantic Conditioned Motion Generation. While previous tasks tackle motion synthesis conditioned by key-poses, other modalities can be used, such as text (Tevet et al. 2022), or 3D scenes (Zheng et al. 2022). Here we focus on the unification of pose-conditioned human motion synthesis, which can serve as a basis to later adopt in other modalities.




Figure 2: UNIMASK-M architecture. Let a human motion \mathbf{X} and its respective binary mask \mathbf{M} . We first interpolate \mathbf{X}_g to obtain \mathbf{X}_{fill} and provide consistency to the input. Then, our Pose Decomposition module (PD) deconstructs each pose \mathbf{p}_t into a sequence of patches $\hat{\mathbf{p}}_t$, which we project and flatten to a sequence of tokens \mathbf{E} . We add the emb_{mix} to \mathbf{E} to inform the encoder and decoder about the masked tokens and the spatio-temporal structure. Our ViT-based encoder and decoder reconstruct the sequence of tokens. Our Pose Aggregation module (PA) regroups the decoded tokens into poses using an MLP layer. Finally, each pose is projected back to the joint representations and summed to our reference motion \mathbf{X}_{ref} .

Masked Autoencoders

Learning robust visual representations is essential for the image domain. With the success of Vision transformers (ViT) (Dosovitskiy et al. 2020), masked visual modeling has emerged as a key design to encode images effectively. Bao, Dong, and Wei (2021) proposed to learn visual representations by predicting the discrete tokens from images and videos. He et al. (2021) proposed an asymmetric encoder-decoder design for masked image modeling, and Tong et al. (2022) adapted this pre-training framework for videos. Recent works have shown the success of Masked Autoencoders (MAE) in reconstructing time-series data (Li et al. 2023), unifying representation learning and image synthesis (Li et al. 2022), or obtaining useful representations for the robotic domain (Karamcheti et al. 2023). In this paper, we propose a novel approach to adapt the image patch decomposition from ViT (Dosovitskiy et al. 2020) to human poses and formulate the motion synthesis as a reconstruction problem inspired by MAEs.

Methodology

In this section, we present our Unified Masked Autoencoder with patchified SKELETONS for Motion synthesis, called UNIMASK-M. First, we reformulate the human motion synthesis task as a reconstruction problem. Then we describe the different components of our approach. A visual illustration of our UNIMASK-M architecture is depicted in Fig. 2.

Problem Formulation. Let $\mathbf{p}_t = [p_{t,1}, \dots, p_{t,J}] \in \mathbb{R}^{J \times n}$ be a human pose at time t composed by J joints and $\mathbf{X} = [\mathbf{p}_0, \dots, \mathbf{p}_T] \in \mathbb{R}^{T \times J \times n}$ a human motion. Each joint $p_{t,j}$ can be described in various representations such as the standard euclidean xyz -position ($n = 3$) or the robust ortho-6D rotation (Zhou et al. 2019) ($n = 6$). Let also a binary mask $\mathbf{M} = [\mathbf{m}_0, \dots, \mathbf{m}_T] \in \mathbb{R}^{T \times J \times n}$ indicate the visibility of each joint in \mathbf{X} . The task of motion synthesis is defined as the reconstruction of the missing joints $\mathbf{X}_m = \mathbf{X} \odot \mathbf{M}$

conditioned by the known motion $\mathbf{X}_g = \mathbf{X} \odot (1 - \mathbf{M})$.

Baseline Strategy. We adapt the input and output motion to facilitate the model’s ability to learn and make accurate predictions, as done in (Valls Mascaro et al. 2022; Oreshkin et al. 2022). First, we fill the missing joints \mathbf{X}_m using interpolation or repeating the last observed information. We denote this filled motion as \mathbf{X}_{fill} , which provides smoothness and continuous flow of motion. By interpolating the masked motion, we ensure no sudden jumps or disruptions in the motion information, which could negatively affect the model’s ability to learn and make accurate predictions. Second, we adopt the delta-strategy from Oreshkin et al. (2022), where the output of our neural network $f_\theta(\mathbf{X}_{fill})$ represents the deviation from an interpolated reference motion \mathbf{X}_{ref} to the ground-truth \mathbf{Y} . Equation 1 summarizes this strategy. This approach has the advantage of handling global reference frame shifts while improving the smoothness of the transitions from the key-poses.

$$\mathbf{Y} = f_\theta(\mathbf{X}_{fill}) + \mathbf{X}_{ref} \quad (1)$$

Pose Decomposition (PD). We propose a novel Pose Decomposition module (PD) to deconstruct a sequence of skeletons (\mathbf{X}_{fill}) into flattened limb-based patches. We group joints belonging to each limb (legs and arms) and trunk, and restructure the human pose \mathbf{p}_t into L patches $\hat{\mathbf{p}}_t = [\hat{\mathbf{p}}_t^0, \dots, \hat{\mathbf{p}}_t^L]$, where each patch $\hat{\mathbf{p}}_t^l$ has a pre-defined number of joints n_l . For a given pose \mathbf{p}_t , each patch $\hat{\mathbf{p}}_t^l$ is projected independently to a token representation $\mathbf{e}_t^l \in \mathbb{R}^D$ through a linear layer. Then, we flatten all patches for the entire motion, creating a sequence of tokens $\mathbf{E} \in \mathbb{R}^{L \cdot T \times D}$. We also restructure the input mask \mathbf{M} into patches, such that $\mathbf{M}_L \in \mathbb{R}^{L \cdot T}$ represents a binary mask per token. For instance, if the binary mask for the token of the trunk has a value of 1, we consider the trunk to be visible.

The PD module serves two purposes. Firstly, it allows the model to condition the motion only on partial information, the tokens, unlike previous approaches that could only




Figure 3: Mixed embedding strategy. The mixed embeddings are obtained by summing (i) a masking token to identify the masked patches (emb_{mask}); (ii) $L = 5$ spatial learnable parameters that correspond to each body part (emb_{kin}); and (iii) a sinusoidal position embedding (emb_{pos}).

consider the whole pose as a condition. This is because we project patches independently to the representation space instead of projecting the full-body pose. Secondly, by flattening the patches into a sequence of tokens, our UNIMASK-M can leverage self-attention to capture both the spatial and temporal relationships in a single step, contrary to the parallel processing used in Valls Mascaro et al. (2022). We demonstrate the performance boost of using PD compared to the existing full-body projection in our ablation study.

Mixed Embeddings. We adopt a mixed embeddings strategy emb_{mix} to inform the model about the spatio-temporal structure of \mathbf{X}_{fill} and the identification of the masked patches \mathbf{M} to be reconstructed, as shown in Fig. 3. For that, we extend the sinusoidal positional encoding (emb_{pos}) used in Transformers with a kinematic embedding (emb_{kin}). emb_{kin} is a patch-dependent learnable parameter added to each token belonging to a particular limb. This embedding informs about the kinematic structure of the input sequence. By incorporating emb_{kin} , the model can distinguish between different body parts and leverage their spatial relationships. Finally, inspired by MAEs, we add a learnable parameter to the missing patches (emb_{mask}).

Masked Autoencoder. We adopt the standard ViT (Dosovitskiy et al. 2020) self-attention encoder and decoder block, composed of Multi-Layer Perceptrons and Layer Normalization (Ba, Kiros, and Hinton 2016). Unlike previous works (He et al. 2021; Tong et al. 2022; Li et al. 2022), we do not neglect the masked patches in the encoder phase, as all patches \mathbf{E} have relevant information for the model (due to the adaptation of \mathbf{X} into \mathbf{X}_{fill}). We verify our hypothesis through an Ablation Study. Finally, both the encoder and the decoder share the same emb_{mix} .

Pose Aggregation (PA). The output of the decoder is a flattened sequence of reconstructed tokens denoted by $\hat{\mathbf{E}} \in \mathbb{R}^{L \cdot T \times D}$. To leverage the hierarchical structure of a human skeleton, we re-group the tokens that correspond to a particular pose into $\hat{\mathbf{E}}_{pose} \in \mathbb{R}^{T \times L \cdot D}$ and then project it to a single pose representation $\mathbf{E}_{pose} \in \mathbb{R}^{T \times D}$ using two fully connected layers with GELU activation (Hendrycks and Gimpel 2016). PA allows fusing the different body-part token representations into a full-body pose, which facilitates




Figure 4: Adaption of UNIMASK-M using DCT and TempMLP. First, we apply Discrete Cosine Transformation (DCT) and Inverted DCT (IDCT) to encode and decode the given motion. Additionally, we adopt a Temporal MLP (TempMLP) module to refine the predicted pose sequence through M blocks of fully connected layers (FC), Layer Normalization (LN) and a residual connection.

the prediction of more plausible motions. Finally, we project \mathbf{E}_{pose} to the joint-based pose $\hat{\mathbf{Y}} \in \mathbb{R}^{T \times P}$ and add it to the reference motion \mathbf{X}_{ref} .

Experimental Results

In this section, we evaluate the effectiveness of our UNIMASK-M model across different tasks and datasets by conducting multiple experiments and comparisons. Unless stated otherwise, all models in comparison are trained for a target task within the proposed dataset. Additional qualitative results can be found in ¹.

Dataset and Metrics. Human motion forecasting has been mainly addressed in Human3.6M dataset (Ionescu et al. 2014). This dataset includes 3.6 million 3D poses of humans performing 15 daily activities. Following standard evaluation, we report Mean Per Joint Position Error (MPJPE) in the test subject S5, as suggested in (Martinez, Black, and Romero 2017; Mao et al. 2019; Mao, Liu, and Salzmann 2020) when predicting the next 1s motion given 400ms of the past. Human motion inbetweening is evaluated in LaFAN1 dataset (Harvey et al. 2020a). This dataset contains 496,672 motions sampled at 30Hz and captured in a MOCAP studio. We report the average L2 distance of the global position (L2P) or rotation (L2Q, in quaternion) per frame, as well as the Normalized Power Spectrum Similarity (NPSS) to evaluate angular differences between prediction and ground truth on the frequency domain. The given key-poses consist of 10 poses from the past and 1 pose from the future as input. The number of intermediate poses to be predicted can vary between 5 and 30, as proposed by (Duan et al. 2021; Oreshkin et al. 2022; Harvey and Pal 2018).

Quantitative Evaluation

Human Motion Forecasting. Our model is tested in Human3.6M dataset and shows competitive results to the state-of-the-art models, as shown in Table 1. Following the standard evaluation in Human3.6M, all depicted models in Table 1 are trained to predict the next 25 poses given 10 poses of

¹<https://evm7.github.io/UNIMASKM-page/>

milliseconds (ms)	80	320	400	560	1000
Repeat last pose	23.8	76.1	88.2	107.4	136.6
(Mao et al. 2019)	12.4	49.9	60.9	79.5	112.7
(Dang et al. 2021)	12.1	51.6	62.9	-	114.2
(Bouazizi et al. 2022)	12.8	53.1	64.5	82.9	117.1
ST-DGCN (Ma et al. 2022)	10.3	46.6	57.5	76.3	110.0
SiMLPe (Guo et al. 2023)	10.4	49.6	61.0	80.3	114.8
EqMotion (Xu et al. 2023)	10.8	51.3	62.8	81.8	115.4
UNIMASK-M (baseline)	17.7	59.8	70.8	89.0	120.6
UNIMASK-M (w/ DCT)	15.8	53.8	64.4	81.7	113.6
UNIMASK-M (Ours)	11.9	50.7	61.6	79.6	112.1

Table 1: MPJPE error in 3D human motion forecasting for Human3.6M dataset when predicting 25 next poses given 10 past poses. Here, bold denotes the best result at each horizon.

milliseconds (ms)	80	320	400	560	1000
(Mao et al. 2019)	32.8	74.4	83.5	97.8	125.8
(Bouazizi et al. 2022)	32.7	76.9	86.8	102.0	129.1
ST-DGCN (Ma et al. 2022)	32.8	74.3	83.5	97.8	124.6
SiMLPe (Guo et al. 2023)	32.7	74.0	82.8	96.7	124.7
UNIMASK-M (Ours)	27.3	64.9	74.5	90.3	120.5

Table 2: MPJPE error in motion forecasting given an observation with 20% joints occluded for the Human3.6M dataset.

the past. In addition to the UNIMASK-M architecture, we adopt the Discrete Cosine Transformation (DCT) to the input motion due to the periodic movements accounted for in this dataset. This strategy is also adopted for most works shown in Table 1 as it boosts the performance (Mao et al. 2019; Ma et al. 2022; Dang et al. 2021; Guo et al. 2023; Xu et al. 2023). Moreover, we smooth the transition from prefix sequence to predicted poses using an additional Temporal MLP (TempMLP) that refines the decoded poses \mathbf{E}_{pose} in time. These modifications, shown in Fig. 4, are adopted only for Human3.6M due to the cyclic motions encountered. We also provide our results without them in Table 1. Contrary to existing autoregressive approaches (Guo et al. 2023; Ma et al. 2022), which excel in the short term but accumulate errors in long-term horizons, our UNIMASK-M model predicts the whole motion in one forward pass. Additionally, auto-regressive models can not handle the in-betweening or completion task, as they only consider the past information as input. A visual comparison is depicted in Figure 5 which shows the accurate prediction of our model in the long term.

On top of that, we also evaluate the motion forecasting performance when the input motion contains occlusions. For that, we randomly mask the joints for the observed sequence with a probability p . For instance, $p = 0.2$ implies that 20% of the joints are not visible. We apply the same approach to pure forecasting models from literature and show their results in Table 2. Results show that our UNIMASK-M can better handle occlusions in the input sequence.

Human Motion Inbetweening. We evaluate our UNIMASK-M model for motion inbetweening in LaFANI dataset, and show the results in Table 3. Qin, Zheng, and Zhou (2022) applies task-specific tricks (foot contact




Figure 5: Comparison of the motion forecasting task. Predicted skeletons are shown in red and blue.

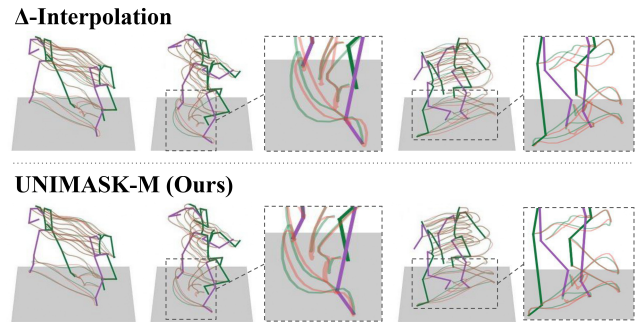


Figure 6: Visual comparison of the inbetweening task with Oreshkin et al. (2022). We show the predicted motion trace for both Oreshkin et al. (2022) (top row) and our UNIMASK-M (bottom row). The results show that our predicted trace (red) is closer to the ground-truth trace (green).

and smooth joint losses, post-processing steps, and trains two models) to boost its performance, but underperforms when using a single fill-in model as ours (ToG_{ctx}). Our UNIMASK-M performs better for longer transitions and with similar results in shorter periods than Oreshkin et al. (2022). Oreshkin et al. (2022) decouples the prediction using cross-attention for the masked frames and self-attention for the observed poses, facilitating the prediction of the transition close to the key-poses. However, our UNIMASK-M provides more robustness in longer transitions and therefore is more appropriate for synthesis tasks. To confirm this claim, we additionally train both Oreshkin et al. (2022) and UNIMASK-M when predicting longer transitions (50 and 70). These results are reported in Table 3 and show how our UNIMASK-M outperforms state-of-the-art for long-term motion inbetweening. Note that these trained long-term models also perform in-betweening for any transition frames. Qualitative results are shown in Fig. 6 and 7.

Human Motion Completion. We assess the motion com-

Number of Transition Frames	L2Q ↓				L2P ↓				NPSS ↓			
	5	15	30	70	5	15	30	70	5	15	30	70
Zero velocity	0.56	1.10	1.51	2.22	1.52	3.69	6.60	8.16	0.0053	0.0522	0.2318	1.2194
SLERP Interpolation	0.22	0.62	0.98	1.72	0.37	1.25	2.32	3.54	0.0023	0.0391	0.2013	1.2385
TG _{rec} (Harvey et al. 2020a)	0.21	0.48	0.83	-	0.32	0.85	1.82	-	0.0025	0.0304	0.1608	-
TG _{complete} (Harvey et al. 2020a)	0.17	0.42	0.69	-	0.23	0.65	1.28	-	0.0020	0.0258	0.1328	-
SSMCT _{local} (Duan et al. 2021)	0.17	0.44	0.71	-	0.23	0.74	1.37	-	0.0019	0.0291	0.1430	-
SSMCT _{global} (Duan et al. 2021)	0.14	0.36	0.61	-	0.22	0.56	1.10	-	0.0016	0.0234	0.1222	-
Δ -Interpolation (Oreshkin et al. 2022)	0.11	0.32	0.57	1.20	0.13	0.47	1.00	1.91	0.0014	0.0217	0.1215	1.0180
ToG _{ctx} (Qin, Zheng, and Zhou 2022)	0.13	0.33	0.60	-	0.17	0.49	1.07	-	0.0015	0.0212	0.1238	-
ToG (Qin, Zheng, and Zhou 2022)	0.10	0.28	0.54	-	0.10	0.39	0.89	-	0.0011	0.0188	0.1124	-
UNIMASK-M (Ours)	0.12	0.32	0.57	1.16	0.14	0.46	0.97	1.80	0.0014	0.0216	0.1204	0.9532

Table 3: Quantitative evaluation of human motion inbetweening on LAFAN1 dataset. A lower score is better. Here, bold indicates the best result. Note that we trained a different model for the 50 and 70 transition frames for (Oreshkin et al. 2022).

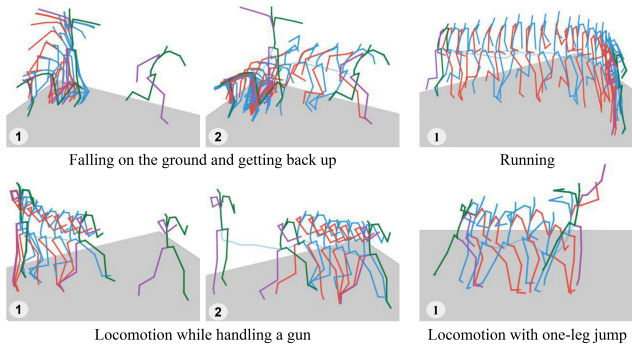


Figure 7: Human motion inbetweening in LaFAN1 Dataset. Given certain key-poses (green and purple skeletons) our UNIMASK-M fills the motion in-between (in red and blue).

pletion task in the Human3.6M dataset by introducing random occlusions throughout the entire motion at varying percentages. Since there is no established benchmark for this topic, we adapt and train Oreshkin et al. (2022) specifically for this purpose by adding the PD module to allow the conditioning based only on body parts. We evaluate the motion completion task when masking 90% of future joints and observe that our UNIMASK-M reduces the MPJPE error by -12.54% compared to this baseline.

Unified Model for Different Tasks. Although our previous results demonstrate the versatility of our model for various tasks, we train a unified model with different masking patterns to achieve high performance in several tasks at once. Additionally, we implement curriculum learning, which gradually increases the probability of masking during training to enhance the model’s robustness in challenging tasks. The advantage of using curriculum learning in our training is shown in Table 4, where the model exhibits performance across tasks comparable to task-specific models.

Comparison With Generative Models. We trained the generative U-CVAE (Cai et al. 2021) and measured the Top1 MPJPE error from 30 synthesized motions. Our UNIMASK-M significantly outperforms U-CVAE in both motion forecasting (+64.49%) and completion (+88.12%). We com-

Motion Forecasting (MPJPE) ↓						
Mask Ratio	80	160	320	400	560	1000
1	11.9	25.1	50.7	61.6	79.6	112.1
0.85 → 1	11.9	24.6	49.9	61.2	80.2	115.7
Motion Completion (MPJPE) ↓						
Mask Ratio	80	160	320	400	560	1000
1	13.0	23.7	41.2	47.5	56.9	82.2
0.9	8.4	14.2	22.4	25.1	29.3	52.0
0.85 → 1	8.5	14.7	22.9	25.7	29.9	51.9

Table 4: MPJPE error of UNIMASK-M in Human3.6M dataset under different training masking probabilities (p_m) and tasks. Training on forecasting requires the whole future mask ($p_m = 1$) while training on completion uses random patch masking ($p_m = 0.9$). To boost performance in both objectives, curriculum learning is applied where p_m increases during training epochs (from $p_m = 0.85$ to $p_m = 1$).

pared qualitatively both in Fig. 5. Motions from U-CVAE are jittery and tend to collapse in the long term. Moreover, U-CVAE is trained only with whole pose masking and generates non-natural behaviors when given partial poses. Finally, MDM (Tevet et al. 2022) does not handle key poses and lacks real-time capabilities. We compute the average run time per 10000 runs and observe that our proposed UNIMASK-M is $494.79\times$ faster than MDM.

Ablation Study

This section analyses the various approaches proposed in this work and their impact on the LaFAN1 dataset. All results are presented in Table 5.

Pose Decomposition (PD). Our UNIMASK-M proposes PD to leverage self-attention among body parts but aggregates the spatial information before the pose projection through PA. We compare our PD approach against the typical linear projection proposed in the literature for pose encoding. We also design an inverted PD that linearly projects each decoded body parts patch \hat{E} to the assigned joints independently, and we compare it with our PA approach. Table 5 shows the performance boost of our design.

Num Frames	L2Q ↓		L2P ↓		NPSS ↓	
	5	30	5	30	5	30
w/o PD	0.13	0.59	0.15	1.03	0.0015	0.1263
w/o PA	0.13	0.60	0.15	1.05	0.0015	0.1271
w/ DCT	0.17	0.68	0.23	1.33	0.0019	0.1503
w/ TempMLP	0.12	0.60	0.14	1.04	0.0014	0.1280
w/o emb_{kin}	0.13	0.58	0.15	0.98	0.0014	0.1216
w/ Attn Mask	0.13	0.59	0.15	1.02	0.0015	0.1243
w/o encoder	0.12	0.58	0.14	0.99	0.0014	0.1248
w/o decoder	0.13	0.60	0.16	1.05	0.0015	0.1282
Light	0.12	0.59	0.14	1.00	0.0014	0.1248
UNIMASK-M	0.12	0.57	0.14	0.97	0.0014	0.1204

Table 5: Results of our UNIMASK-M under different configurations in the inbetweening task and the LaFan1 dataset.

DCT and TempMLP. We evaluate our model with DCT and Temporal MLPs, as depicted in Fig. 4, and observe that DCT highly degrades the model’s performance. LaFAN1 dataset contains more complex and less periodical motions, therefore working on the frequency domain results in an information loss. On the other hand, plugging an additional Temporal MLP causes the model to overfit.

Attention Mechanism. To rely on the self-attention for both masked and key-poses, our model takes advantage of the mixed embeddings to be informed about (i) the temporal relations, (ii) the masked poses, and (iii) the spatial relations. While (i) and (ii) are already proposed for motion synthesis (Duan et al. 2021), (iii) is considered a novelty due to the uniqueness of the PD. We show the importance of using emb_{kin} in Table 5. We also train our model with a triangular attention mask to enforce causality. As expected due to the bi-directionality of the inbetweening tasks, imposing causal attention degrades performance.

Model Size. The final UNIMASK-M proposed for motion inbetweening in LaFAN1 dataset has 41.4M parameters, with 3M parameters less than Oreshkin et al. (2022). We evaluate our model using different configurations of depth and width in the encoder-decoder structure. We observe that changing the depth of the model does not have a high impact on the performance while decreasing the width has a direct effect on the motion quality for longer transitions. Still, our *UNIMASK-M light* performs better than most previous approaches while having only 12.4M parameters. We also remove the encoder or decoder to reduce the model size, which causes performance degradation. Our UNIMASK-M still exhibits high performance only when using the decoder.

Effect of Masking Ratio. We analyze different masking ratios for our UNIMASK-M in motion forecasting with occlusions (see Table. 2). A masking ratio over 50% increases the MPJPE short-term error by 2.65, but has a marginal impact in the long term ($\times 1.05$). These findings suggest a strong correlation between the short-term poses and the observed motion, highlighting the vulnerability of auto-regressive approaches (Guo et al. 2023; Ma et al. 2022) to occlusions, where their MPJPE error increases by 1.19 and 1.14 with 40% occlusion respect to UNIMASK-M.

Strategy	80	320	400	560	1000
S1 (joints)	12.44	51.52	62.59	81.02	114.06
S2 (w/o PD)	12.38	51.57	62.62	80.60	113.05
S3 (ours)	11.85	50.69	61.57	79.58	112.10
S4	11.79	51.08	62.28	80.87	113.62
S5	11.53	50.24	61.38	79.86	113.38

Figure 8: MPJPE millimeter error of the motion forecasting task of UNIMASK-M under different patch granularity.

Curriculum Learning vs Pretraining. Some prior Masked-Autoencoders (MAEs) works benefit from an initial pre-training stage, and later finetuning for downstream tasks. We tested this approach and pretrained our UNIMASK-M using random masks for different motion synthesis tasks and later finetuned for motion completion. The results show a higher MPJPE error (+0.89 on average) than the curriculum learning shown in Table 4.

Granularity of the Patch Hierarchy. We trained different patch granularities and show the most relevant in Fig. 8 for motion forecasting. The results indicate a superior performance on average of our proposed decomposition [S3], with better performance in long-term forecasting. Extremely high [S1] and low [S2] granularities underperform, reinforcing the motivation behind our work. Additionally, considering the clavicle as part of the trunk [S4] or the hips as a separate patch [S5] improves short-term predictions, although the motion collapses in the long term. Visually, having a larger trunk [S2, S4, S5] produces slower movements. In contrast, higher granularity [S1] produces longer sequences of patches, thus increasing the inference time ($\times 1.19\%$) and hindering the model’s ability to capture the spatiotemporal relationships, often resulting in a collapse to a certain pose. These behaviors are depicted in Fig. 5 for UNIMASK-M (w/o PD) and UNIMASK-M (joints).

Conclusion

In this paper, we proposed a unified Masked Autoencoder with patchified SKELETONS for Motion synthesis (UNIMASK-M). By considering all motion synthesis tasks as a masked reconstruction problem, we design a task-agnostic model that can achieve competitive performance compared to the task-dependent state-of-the-art. Ours is the first work to adopt masked autoencoders for general motion synthesis, which might serve as a basis for future works in the field. In addition, we show how decomposing a human pose into patches provides not only more flexible pose conditioning but also boosts performance. We also demonstrate the efficiency of UNIMASK-M design for synthesizing robust motions in occluded or partial conditions. In conclusion, our approach demonstrates its effectiveness and versatility in generating high-quality human motion synthesis.

Acknowledgements

This work is funded by Marie Skłodowska-Curie Action Horizon 2020 (Grant agreement No. 955778) for project 'Personalized Robotics as Service Oriented Applications' (PERSEO).

References

- Ahn, H.; Valls Mascaro, E.; and Lee, D. 2023. Can We Use Diffusion Probabilistic Models for 3D Motion Prediction? In *2023 IEEE International Conference on Robotics and Automation (ICRA)*.
- Aksan, E.; Kaufmann, M.; Cao, P.; and Hilliges, O. 2021. A spatio-temporal transformer for 3d human motion prediction. In *International Conference on 3D Vision (3DV)*, 565–574. IEEE.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bao, H.; Dong, L.; and Wei, F. 2021. BEiT: BERT Pre-Training of Image Transformers. *arXiv preprint arXiv:2106.08254*.
- Baradel, F.; Brégier, R.; Groueix, T.; Weinzaepfel, P.; Kalantidis, Y.; and Rogez, G. 2022. PoseBERT: A Generic Transformer Module for Temporal 3D Human Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bouazizi, A.; Holzbock, A.; Kressel, U.; Dietmayer, K.; and Belagiannis, V. 2022. MotionMixer: MLP-based 3D Human Body Pose Forecasting. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 791–798. International Joint Conferences on Artificial Intelligence Organization.
- Bütpage, J.; Kjellström, H.; and Kragic, D. 2018. Anticipating many futures: Online human motion prediction and generation for human-robot interaction. In *International Conference on Robotics and Automation (ICRA)*, 4563–4570. IEEE.
- Cai, Y.; Wang, Y.; Zhu, Y.; Cham, T.-J.; Cai, J.; Yuan, J.; Liu, J.; Zheng, C.; Yan, S.; Ding, H.; Shen, X.; Liu, D.; and Thalmann, N. M. 2021. A Unified 3D Human Motion Synthesis Model via Conditional Variational Auto-Encoder. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11645–11655.
- Dang, L.; Nie, Y.; Long, C.; Zhang, Q.; and Li, G. 2021. MSR-GCN: Multi-Scale Residual Graph Convolution Networks for Human Motion Prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11467–11476.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Duan, Y.; Shi, T.; Zou, Z.; Lin, Y.; Qian, Z.; Zhang, B.; Lab, Y. Y. N. F. A.; of Michigan, U.; and NetEase. 2021. Single-Shot Motion Completion with Transformer. *ArXiv*, abs/2103.00776.
- Fragkiadaki, K.; Levine, S.; Felsen, P.; and Malik, J. 2015. Recurrent network models for human dynamics. In *Proceedings of the IEEE international conference on computer vision*, 4346–4354.
- Guo, W.; Du, Y.; Shen, X.; Lepetit, V.; Alameda-Pineda, X.; and Moreno-Noguer, F. 2023. Back to MLP: A Simple Baseline for Human Motion Prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 4809–4819.
- Harvey, F. G.; and Pal, C. J. 2018. Recurrent transition networks for character locomotion. *SIGGRAPH Asia 2018 Technical Briefs*.
- Harvey, F. G.; Yurick, M.; Nowrouzezahrai, D.; and Pal, C. 2020a. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4): 60–1.
- Harvey, F. G.; Yurick, M.; Nowrouzezahrai, D.; and Pal, C. J. 2020b. Robust motion in-betweening. *ACM Trans. Graph.*
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2021. Masked Autoencoders Are Scalable Vision Learners.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 6840–6851.
- Holden, D.; Saito, J.; and Komura, T. 2016. A Deep Learning Framework for Character Motion Synthesis and Editing. *ACM Trans. Graph.*
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(7): 1325–1339.
- Jain, A.; Zamir, A. R.; Savarese, S.; and Saxena, A. 2016. Structural-rnn: Deep learning on spatio-temporal graphs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 5308–5317.
- Jiang, J.; Chen, J.; and Guo, Y. 2022. A Dual-Masked Auto-Encoder for Robust Motion Capture with Spatial-Temporal Skeletal Token Completion. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM '22, 5123–5131. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392037.
- Karamcheti, S.; Nair, S.; Chen, A. S.; Kollar, T.; Finn, C.; Sadigh, D.; and Liang, P. 2023. Language-Driven Representation Learning for Robotics.
- Kim, J.; Byun, T.; Shin, S.; Won, J.; and Choi, S. 2022. Conditional motion in-betweening. *Pattern Recognition*, 132.
- Kingma, D. P.; Welling, M.; et al. 2019. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4): 307–392.
- Li, T.; Chang, H.; Mishra, S. K.; Zhang, H.; Katabi, D.; and Krishnan, D. 2022. Mage: Masked generative encoder to unify representation learning and image synthesis. *arXiv preprint arXiv:2211.09117*.
- Li, Z.; Rao, Z.; Pan, L.; Wang, P.; and Xu, Z. 2023. Ti-MAE: Self-Supervised Masked Time Series Autoencoders.

- Ma, T.; Nie, Y.; Long, C.; Zhang, Q.; and Li, G. 2022. Progressively Generating Better Initial Guesses Towards Next Stages for High-Quality Human Motion Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6437–6446.
- Mao, W.; Liu, M.; and Salzmann, M. 2020. History repeats itself: Human motion prediction via motion attention. In *European Conference on Computer Vision (ECCV)*, 474–489. Springer.
- Mao, W.; Liu, M.; Salzmann, M.; and Li, H. 2019. Learning trajectory dependencies for human motion prediction. In *International Conference on Computer Vision (ICCV)*, 9489–9497.
- Martinez, J.; Black, M. J.; and Romero, J. 2017. On human motion prediction using recurrent neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2891–2900.
- Oreshkin, B. N.; Valkanas, A.; Harvey, F. G.; Ménard, L.-S.; Bocquet, F.; and Coates, M. J. 2022. Motion Inbetweening via Deep Δ -Interpolator. arXiv:2201.06701.
- Qin, J.; Zheng, Y.; and Zhou, K. 2022. Motion In-Betweening via Two-Stage Transformers. *ACM Trans. Graph.*, 41(6).
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Bermano, A. H.; and Cohen-Or, D. 2022. Human Motion Diffusion Model. *arXiv preprint arXiv:2209.14916*.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Video-MAE: Masked Autoencoders are Data-Efficient Learners for Self-Supervised Video Pre-Training. In *Advances in Neural Information Processing Systems*.
- Valls Mascaro, E.; Ma, S.; Ahn, H.; and Lee, D. 2022. Robust human motion forecasting using transformer-based model. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10674–10680.
- 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 (NeurIPS)*, 30.
- Xu, C.; Tan, R. T.; Tan, Y.; Chen, S.; Wang, Y. G.; Wang, X.; and Wang, Y. 2023. EqMotion: Equivariant Multi-agent Motion Prediction with Invariant Interaction Reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1410–1420.
- Zheng, Y.; Yang, Y.; Mo, K.; Li, J.; Yu, T.; Liu, Y.; Liu, C. K.; and Guibas, L. J. 2022. Gimo: Gaze-informed human motion prediction in context. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, 676–694. Springer.
- Zhou, Y.; Barnes, C.; Lu, J.; Yang, J.; and Li, H. 2019. On the Continuity of Rotation Representations in Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhu, W.; Ma, X.; Liu, Z.; Liu, L.; Wu, W.; and Wang, Y. 2022. Learning Human Motion Representations: A Unified Perspective. *arXiv preprint arXiv:2210.06551*.