

# 000 SIMPLIHUMON: SIMPLIFYING HUMAN MOTION PRE- 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DICTION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Human motion prediction combines the tasks of trajectory forecasting, human pose prediction, and possibly also multi-person modeling. For each of the three tasks, specialized, sophisticated models have been developed due to the complexity and uncertainty of human motion. While compelling for each task, combining these models for holistic human motion prediction is non-trivial. Conversely, holistic human motion prediction methods, which have been introduced recently, have struggled to compete on established benchmarks for individual tasks. To address this dichotomy, we study a simple yet effective model for human motion prediction based on a transformer architecture. The model employs a stack of self-attention modules to effectively capture both spatial dependencies within a pose and temporal relationships across a motion sequence. This simple, streamlined, end-to-end model is sufficiently versatile to handle pose-only, trajectory-only, and combined prediction tasks without task-specific modifications. We demonstrate that our approach achieves state-of-the-art results across all tasks through extensive experiments on a wide range of benchmark datasets, including Human3.6M, AMASS, ETH-UZY, and 3DPW. Our results challenge the prevailing notion that architectural complexity is a prerequisite for achieving accuracy and generality in human motion prediction. Code will be released.

## 1 INTRODUCTION

Human motion prediction, the task of forecasting future 3D human motion from a sequence of past observations, is a critical challenge with wide-ranging applications in autonomous driving (Zheng et al., 2022; Paden et al., 2016), robotics (Zou, 2024; Salzmann et al., 2023), virtual reality (Clark et al., 2020; Fu et al., 2020; Ro et al., 2019), and sports analytics (Li et al., 2021). Because human motion is inherently multi-dimensional, non-linear, and highly uncertain, the literature has largely tackled prediction of human motion by addressing distinct tasks individually: trajectory prediction (Gu et al., 2022; Bae et al., 2022; Shi et al., 2023; Bae et al., 2024; Yao et al., 2024; Fang et al., 2025), pose prediction (Dang et al., 2022; Barquero et al., 2023; Sun & Chowdhary, 2024; Hosseininejad et al., 2025; Curreli et al., 2025; Xu et al., 2024), and multi-person motion prediction (Jeong et al., 2024; Zheng et al., 2025).

While making individual tasks easier to address, this differentiation also opens up a gap: tasks like pose and trajectory forecasting are fundamentally interrelated and governed by the same underlying dynamics (Zheng et al., 2025), yet they are modeled separately using task-specific architectures. This has led to the development of complex, specialized models that excel at one task but struggle to generalize, limiting their applicability and introducing unnecessary complexity. Notable exceptions that jointly model these different tasks, particularly in the context of multi-person motion, are Jeong et al. (2024) and Zheng et al. (2025). However, the results of these holistic models are suboptimal on established benchmarks for individual sub-tasks. Consequently, models that predict jointly tend to create their own benchmarks or evaluation protocols, making it difficult to assess their effectiveness against specialized methods directly. Their performance limitations on pose and trajectory prediction show the need for a solution that not only addresses human motion prediction holistically but also excels on established, task-specific benchmarks.

To achieve this, we present a general and, in hindsight, very simple approach to 3D human motion prediction. Our model is built upon a stack of self-attention modules to effectively capture

054 both the spatial dependencies within a single pose and the temporal relationships across the entire  
 055 motion sequence. This design allows us to model a variety of complex motion dynamics while main-  
 056 taining a streamlined and efficient framework. Unlike more complicated, multi-stage models, our  
 057 method employs a unified, end-to-end training process, which improves training stability and over-  
 058 all performance. Our findings demonstrate that a well-designed, attention-based model can achieve  
 059 benchmark performance across all tasks, challenging the notion that architectural complexity is a  
 060 prerequisite for accuracy and generality in this field.

061 We validate our approach through extensive experiments on a wide range of public datasets, includ-  
 062 ing Human3.6M (Ionescu et al., 2013) and AMASS (Mahmood et al., 2019) for pose prediction,  
 063 ETH-UCY (Lerner et al., 2007; Pellegrini et al., 2009) and SDD (Robicquet et al., 2016) for tra-  
 064 jectory prediction, as well as MOCAP-UMPM (CMU Graphics Lab, 2003; van der Aa et al., 2011)  
 065 and 3DPW (von Marcard et al., 2018) for combined pose and trajectory tasks. Our results show  
 066 that our model outperforms or matches current best methods across various metrics while being  
 067 computationally efficient.

068 The key contributions of this paper are summarized as follows:  
 069

- 070 • We introduce SimpliHuMoN, a **simple, unified transformer architecture that can outper-**  
 071 **form results of complex, specialized human motion prediction models**  
**unified Transformer**  
**framework that challenges the prevailing trend of architectural complexity in human motion**  
**prediction.**
- 074 • We establish state-of-the-art performance across pose, trajectory, and holistic prediction  
 075 tasks **, showing that a single, simple architecture can outperform highly specialized models.**

## 077 2 SIMPLIHUMON

079 We propose a simple yet effective 3D human motion prediction model based on a transformer de-  
 080 coder architecture. The model is designed to be as simple as possible, learning a mapping from a  
 081 person’s past movements to their future movements while accommodating various input and output  
 082 configurations.

083 The input  $X_{\text{past}}$  consists of two components, each over a historical time horizon of  $H$  timesteps. On  
 084 the one hand, the trajectory  $T_{\text{past}} \in \mathbb{R}^{H \times 3}$  represents the path of a root joint (e.g., the hip). On the  
 085 other hand, the relative body pose  $P_{\text{past}} \in \mathbb{R}^{H \times M \times 3}$  represents the state of  $M$  joints relative to the  
 086 root joint. Our framework can operate on either of these inputs individually or on both combined: for  
 087 trajectory prediction, the model only operates on  $T_{\text{past}}$ ; for pose prediction, the model only operates  
 088 on  $P_{\text{past}}$ ; and for joint pose and trajectory prediction, the model operates on both.

089 The model aims to predict the corresponding future state  $X_{\text{fut}}$ , over a prediction horizon of  $F$   
 090 timesteps. To capture the uncertainty of motion, following prior work (Jeong et al., 2024), the model  
 091 generates  $K$  distinct proposal states, *i.e.*,  $X_{\text{fut}} = (X_{\text{fut}}^1, \dots, X_{\text{fut}}^K)$ . Each proposal  $X_{\text{fut}}^k$ ,  $k \in \{1, \dots, K\}$ ,  
 092 consists of a complete predicted future state. The composition of  $X_{\text{fut}}^k$  mirrors that of the input; it  
 093 can include a future root trajectory  $T_{\text{fut}} \in \mathbb{R}^{F \times 3}$ , a future relative body pose  $P_{\text{fut}} \in \mathbb{R}^{F \times M \times 3}$ , or  
 094 both, depending on what was provided as input.

096 **Overview of our method.** As illustrated in Fig. 1, our model begins by independently processing  
 097 the historical observations  $X_{\text{past}}$  and a set of learnable query tokens  $\mathcal{Q}_{\text{in}} = (\mathcal{Q}_{\text{in}}^1, \dots, \mathcal{Q}_{\text{in}}^F) \in \mathbb{R}^{F \times 3}$   
 098 into a context tensor  $\mathcal{C}$  and a query tensor  $\mathcal{Q}$  respectively (Sec. 2.1). A self-attention-based trans-  
 099 former then processes the tensors (Sec. 2.2). Finally, a multi-modal prediction head regresses the  
 100 decoder’s output  $Z$  into  $K$  distinct trajectories and pose hypotheses to give the final output,  $X_{\text{fut}}$   
 101 (Sec. 2.3). We describe the training procedure and model configurations in Sec. 2.4.

### 103 2.1 INPUT PROCESSING AND EMBEDDING MODULE

105 This module prepares the raw input data for the transformer decoder by normalizing it and map-  
 106 ping it into a shared high-dimensional latent space of dimension  $d_{\text{model}}$ . The process creates two  
 107 main tensors: a context tensor,  $\mathcal{C}$ , from historical observations and a query tensor,  $\mathcal{Q}$ , from a set of  
 learnable parameters.

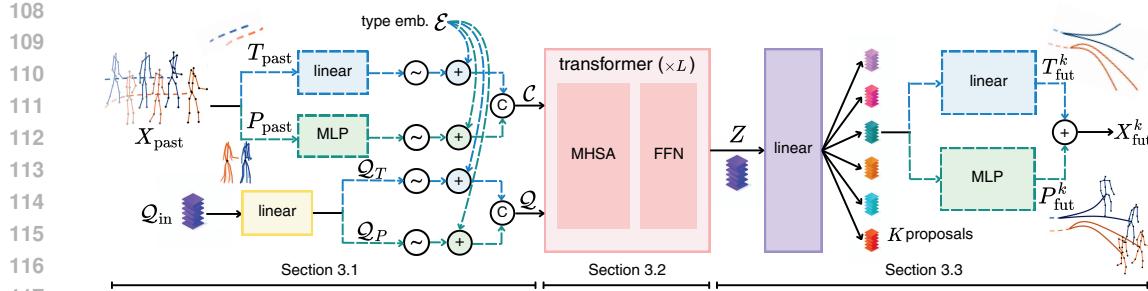


Figure 1: An overview of our architecture. Past observations of 3D poses ( $P_{\text{past}}$ ) and trajectories ( $T_{\text{past}}$ ) are jointly processed by an encoder. Learnable input queries ( $Q_{\text{in}}$ ), representing potential future states, interact with the encoded past motion within a decoder to produce  $K$  distinct future motion proposals ( $X_{\text{fut}}^k$ ) for all agents over a specified horizon.

### 2.1.1 PAST CONTEXT ENCODING

To compute the context tensor  $\mathcal{C}$ , the historical input sequence is processed using one or both of two parallel streams, depending on the task: one for trajectory  $T_{\text{past}}$  and one for relative body pose  $P_{\text{past}}$ .

**Root Trajectory Processing.** The 3D coordinates of the root joint are extracted from the input sequence. To normalize the motion, the root’s position at the final input frame is subtracted from all historical root positions. This normalized trajectory is then projected into the  $d_{\text{model}}$ -dimensional embedding space by a linear layer.

**Relative Pose Processing.** The pose is represented relative to the root (hip) joint for each timestep. If a dataset provides absolute coordinates, we normalize the pose by subtracting the root joint’s position from all other body joint positions. This relative pose vector is then processed by a two-layer MLP (with a GELU activation function), which outputs an embedding of dimension  $d_{\text{model}}$ .

After their initial embedding, both streams are enhanced. First, a sinusoidal positional encoding is added to each sequence to encode the specific position of each of the  $H$  timesteps along the time axis. Then, a learnable type embedding  $\mathcal{E}$  is added to each token. The type embedding encodes whether a given token represents part of the root trajectory or the body pose. Finally, the processed sequences are concatenated (if both are present) along the sequence dimension to form the final context tensor,  $\mathcal{C}$ . The shape of  $\mathcal{C}$  is therefore  $\mathbb{R}^{2H \times d_{\text{model}}}$  for combined inputs, and  $\mathbb{R}^{H \times d_{\text{model}}}$  when only a single input modality is provided.

### 2.1.2 FUTURE QUERY GENERATION

The queries used to prompt the decoder are learnable tensors  $Q_{\text{in}} \in \mathbb{R}^{F \times 3}$ . Similar to the object queries in DETR (Carion et al., 2020) or learnable soft prompts in language modeling (Lester et al., 2021), these are input-independent parameters optimized during training. They serve as initial “slots” for the  $F$  future timesteps, providing the decoder with a temporal structure to fill based on the context. These learnable prompts guide the decoder in its computation. These tokens are first projected into the  $d_{\text{model}}$  space by a linear layer. The resulting sequence is then explicitly split into trajectory  $Q_T \in \mathbb{R}^{F \times d_{\text{model}}}$  and pose  $Q_P \in \mathbb{R}^{F \times d_{\text{model}}}$  queries if both modalities are required. Similar to the past context encoding, these query sequences are enriched with positional encodings and their corresponding type embeddings. The two query sequences are then concatenated (if both are present) to create the final query tensor,  $Q$  ( $\in \mathbb{R}^{2F \times d_{\text{model}}}$  for combined inputs,  $\mathbb{R}^{F \times d_{\text{model}}}$  for single), ensuring that it perfectly mirrors the composition and format of  $\mathcal{C}$ .

This explicit separation of queries into trajectory and pose streams enables the model’s flexibility. The architecture learns a strong association between each query type and its corresponding output modality, reinforced by the type embeddings. This allows the same model to handle different tasks without any architectural modifications.

162  
163

## 2.2 TRANSFORMER DECODER

164  
165  
166  
167

The major computations in our model are performed by a decoder-only transformer with  $L$  identical layers, utilizing a pre-LayerNorm configuration. Each layer operates on the concatenation of two input tensors: a context tensor,  $\mathcal{C}$ , derived from historical observations, and a query tensor,  $\mathcal{Q}$ , derived from learnable latent variables.

168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181

The core of SimpliHuMoN is a decoder-only transformer that processes historical context and future queries as a single, continuous sequence. Unlike standard encoder-decoder architectures that separate inputs into distinct processing streams connected only by cross-attention, we concatenate the context  $\mathcal{C}$  and query  $\mathcal{Q}$  tensors along the temporal dimension to form a unified input sequence  $[\mathcal{C}; \mathcal{Q}] \in \mathbb{R}^{(H+F) \times d_{\text{model}}}$  for self-attention. A key distinction from standard encoder-decoder or cross-attention-based models is our use of a unified attention mechanism. Within each layer, we perform a single multi-head self-attention operation over the sequence  $[\mathcal{C}; \mathcal{Q}]$  concatenated over the time dimension. This design allows every token in the context and query sequences to directly attend to all other tokens, providing a global exchange of information in a single step. For enhanced training stability, we employ pre-LayerNorm with Root Mean Square Layer Normalization (RMSNorm) for training stability and a standard Feed-Forward Network (FFN) with GELU activation. apply Root Mean Square Layer Normalization (RMSNorm) to the query and key projections within each attention head before the dot-product operation. The standard feed-forward network (FFN) sub-layer uses a GELU activation.

182  
183  
184  
185

After passing through the stack of  $L$  decoder layers, the model produces an output tensor,  $Z$ , with the exact dimensions as the input query  $\mathcal{Q}$ . Having attended to the full context, these output query tokens now serve as rich, context-aware representations ready to be mapped into future predictions by the output heads.

186  
187  
188  
189  
190  
191  
192  
193

This unified architecture is task-agnostic: whether the input  $\mathcal{C}$  contains only trajectory, only pose, or both, the self-attention mechanism naturally adapts to model the available dependencies. The decoder's ability to handle different prediction tasks is a direct consequence of this unified attention design. The architecture is agnostic to the composition of the context  $\mathcal{C}$ . For combined prediction, the trajectory and pose queries can attend to their corresponding context streams. If the task is trajectory-only,  $\mathcal{C}$  will only contain trajectory information, and the query tokens  $\mathcal{Q}$  will attend to this relevant context. This allows the model to implicitly specialize its query representations based on the available input, providing a flexible foundation for all task variations.

194  
195

## 2.3 MULTI-MODAL PREDICTION HEADS

196  
197  
198  
199  
200  
201  
202  
203  
204  
205

To account for the stochastic nature of the prediction task, the prediction head decodes the final latent representation from the decoder into  $K$  distinct future hypotheses. The mechanism is a single linear projection from the decoder's output tensor  $Z$  (shape  $[F, d_{\text{model}}]$ ) to an output tensor of shape  $[F, K \times C]$ , where  $C$  is the output dimension (e.g. 3 for trajectory,  $M \times 3$  for pose). This is then reshaped to  $[F, K, C]$ , creating  $K$  parallel branches. Two dedicated output heads then process each branch, if both are being modeled, to regress the future root trajectory ( $T_{\text{fut}}^k$ ) and body pose ( $P_{\text{fut}}^k$ ), respectively, ensuring each of the  $K$  proposals is a complete and comparable hypothesis. Architecturally, these heads mirror the input processing module: a linear layer regresses the trajectory and a two-layer MLP regresses the pose, effectively inverting the initial embedding process.

206  
207

## 2.4 IMPLEMENTATION DETAILS

208  
209  
210  
211

The model is trained end-to-end using a “winner-takes-all” loss, where gradients are backpropagated only through the single hypothesis  $k$  that minimizes the Euclidean distance to the ground truth future. Formally, the training loss  $\mathcal{L}$  for a given ground truth  $X_{\text{fut}}^{\text{gt}}$  is computed via

212  
213

$$\mathcal{L}(X_{\text{past}}, X_{\text{fut}}^{\text{gt}}) = \min_{k \in \{1, \dots, K\}} \|X_{\text{fut}}^{\text{gt}} - X_{\text{fut}}^k(X_{\text{past}})\|_2, \quad (1)$$

214  
215

where  $X_{\text{fut}}^k(X_{\text{past}})$  is the  $k^{\text{th}}$  prediction hypothesis computed from  $X_{\text{past}}$  via the model. This formulation ensures that gradients are only computed for the best prediction, encouraging the model's  $K$  output modes to specialize and cover diverse, plausible futures.

216 We report results for two configurations: a “wide” model ( $L = 6, d_{\text{model}} = 192$ ) and a “deep” model  
 217 ( $L = 16, d_{\text{model}} = 48$ ). In all experiments, we use the AdamW optimizer ( $\beta_1 = 0.95, \beta_2 = 0.999$ )  
 218 with a weight decay of  $10^{-4}$ . All models are trained for 300 epochs with a batch size of 64 and  
 219 standard data augmentation on one NVIDIA RTX A6000 GPU. The number of modes,  $K$ , is set as  
 220 a hyperparameter to follow prior work per task.  
 221

### 222 3 EXPERIMENTS

#### 224 3.1 DATASETS

226 We evaluate our model on several standard benchmarks to cover a range of motion forecasting tasks.  
 227 For 3D human pose prediction, we use Human3.6M (Ionescu et al., 2013), a large-scale lab-based  
 228 dataset, and AMASS (Mahmood et al., 2019), a comprehensive motion capture archive used for  
 229 generative modeling. For trajectory forecasting, we use the pedestrian datasets ETH-UCY (Lerner  
 230 et al., 2007; Pellegrini et al., 2009) and the Stanford Drone Dataset (SDD) (Ro et al., 2019), which  
 231 contains varied persons from an aerial view. Finally, we evaluate joint pose and trajectory prediction  
 232 using Mocap-UMPM (CMU Graphics Lab, 2003; van der Aa et al., 2011), a mixed dataset of  
 233 Mocap and UMPM containing synthesized human interaction between three people, and 3DPW  
 234 (von Marcard et al., 2018), a dataset with two people traversing a real-world environment. We report  
 235 results on each benchmark after training our model on its respective dataset in Table 1, which uses  
 236 the same color scheme to visually group the results by task.  
 237

#### 239 3.2 METRICS

241 We evaluate our model following common practice for multi-modal models that generate  $K$  pro-  
 242 posals, reporting the minimum error among all generated proposals. For pose prediction, we report  
 243 the minimum Average/Final Displacement Error (ADE/FDE) averaged across all body joints over  
 244  $K = 7$  proposals, following Hosseininejad et al. (2025). For trajectory prediction, we report the  
 245 ADE/FDE on the root joint over  $K = 20$  proposals, following Yao et al. (2024). In the combined  
 246 pose and trajectory prediction task, we assess local and global accuracy over  $K = 6$  proposals,  
 247 following Jeong et al. (2024). For this, we use two metrics: Aligned mean per joint Position Error  
 248 (APE), which measures pose error after root-alignment, and Joint Precision Error (JPE), which mea-  
 249 sures the overall error of all joints in the world coordinate system. Consistent with prior work, for  
 250 datasets containing multiple people, the final reported metric is the average of the errors computed  
 251 for each individual.

#### 253 3.3 BASELINES

255 We compare our method against a wide range of state-of-the-art models across three distinct pre-  
 256 diction tasks. In the domain of pose-only prediction, we evaluate against several recent generative  
 257 approaches, including DivSamp (Dang et al., 2022), and prominent diffusion-based models such as  
 258 BeLFusion (Barquero et al., 2023), CoMusion (Sun & Chowdhary, 2024), and SkeletonDiff (Cur-  
 259 reli et al., 2025). Our comparison in this category also includes Motionmap (Hosseininejad et al.,  
 260 2025) and the state-space diffusion model SLD (Xu et al., 2024). For trajectory-only prediction, we  
 261 benchmark against MID (Gu et al., 2022), GP-Graph (Bae et al., 2022), TUTR (Shi et al., 2023),  
 262 SingularTrajectory (Bae et al., 2024), the vision-language model TrajCLIP (Yao et al., 2024), and  
 263 NMRF (Fang et al., 2025). Finally, for the comprehensive task of multi-person motion prediction,  
 264 which involves forecasting combined human trajectory and pose, we include EMPMP (Zheng et al.,  
 265 2025) and T2P (Jeong et al., 2024).

#### 266 3.4 QUANTITATIVE RESULTS

268 Our proposed simple model demonstrates versatile and robust performance, improving state-of-  
 269 the-art results across a diverse range of motion forecasting tasks, as shown in Table 1. Its suc-  
 270 cess as a generalist architecture is particularly noteworthy given that many competing methods are

270 Table 1: Detailed comparison of model performance. Lower values are better ( $\downarrow$ ), with the best  
 271 results shown in **bold**. An asterisk (\*) denotes models we recomputed for this setup, a dagger ( $\dagger$ )  
 272 marks models adapted for the specific task, while a ( $\wedge$ ) notes models that use external training data.  
 273

274	275	276	277		278		279		280		281											
			282		283		284		285		286											
			288		289		290		291		292											
293	294	295	296	297	298	299	300	301	302	303	304	305										
296			297			298			299			300										
297			298			299			300			301										
298			299			300			301			302										
299			300			301			302			303										
300			301			302			303			304										
301			302			303			304			305										
302			303			304			305			306										
303			304			305			306			307										
304			305			306			307			308										
305			306			307			308			309										
306			307			308			309			310										
307			308			309			310			311										
308			309			310			311			312										
309			310			311			312			313										
310			311			312			313			314										
311			312			313			314			315										
312			313			314			315			316										
313			314			315			316			317										
314			315			316			317			318										
315			316			317			318			319										
316			317			318			319			320										
317			318			319			320			321										
318			319			320			321			322										
319			320			321			322			323										
320			321			322			323			324										
321			322			323			324			325										
322			323			324			325			326										
323			324			325			326			327										
324			325			326			327			328										
325			326			327			328			329										
326			327			328			329			330										
327			328			329			330			331										
328			329			330			331			332										
329			330			331			332			333										
330			331			332			333			334										
331			332			333			334			335										
332			333			334			335			336										
333			334			335			336			337										
334			335			336			337			338										
335			336			337			338			339										
336			337			338			339			340										
337			338			339			340			341										
338			339			340			341			342										
339			340			341			342			343										
340			341			342			343			344										
341			342			343			344			345										
342			343			344			345			346										
343			344			345			346			347										
344			345			346			347			348										
345			346			347			348			349										
346			347			348			349			350										
347			348			349			350			351										
348			349			350			351			352										
349			350			351			352			353										
350			351																			

324 but also reveals a clear opportunity for future work: integrating an explicit interaction mechanism  
 325 could yield even better results.  
 326

327 Additionally, we want to note that  
 328 our model is computationally very  
 329 efficient. To demonstrate this,  
 330 we benchmarked all models that  
 331 perform joint pose and trajectory  
 332 prediction on the MOCAP-UMPM  
 333 dataset, comparing the average num-  
 334 ber of samples processed per sec-  
 335 ond. Our “deep” configuration is not  
 336 only more accurate but also more  
 337 computationally efficient than the  
 338 lightweight EMPMP model, show-  
 339 ing a 14.3% increase in training  
 340 throughput and processing test sam-  
 341 ples nearly 1.8 times faster. Please see Table 2 for details.  
 342

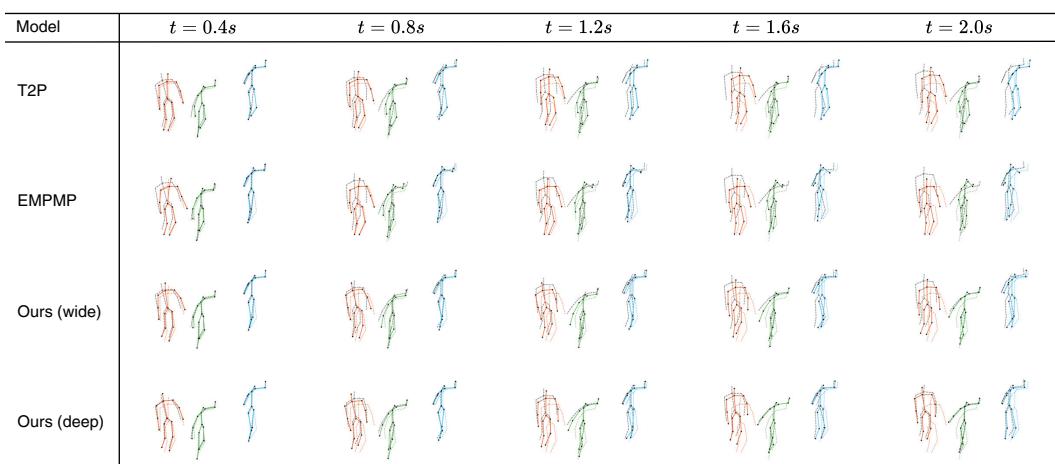
### 343 3.5 QUALITATIVE RESULTS

344 We provide a qualitative comparison of predicted motions on the MOCAP-UMPM dataset in Fig-  
 345 ure 2. The figure illustrates a challenging sample where three individuals are walking backward,  
 346 a motion that requires complex coordination. Our “wide” and “deep” models both generate fluid  
 347 and physically plausible motion sequences that accurately capture the underlying dynamics. The  
 348 articulation of the arms and torso is notably realistic, showcasing the model’s ability to learn natu-  
 349 ral human motion without being constrained by explicit structural priors. In particular, our “deep”  
 350 configuration demonstrates exceptional performance over the long term, maintaining high-quality,  
 351 dynamic predictions even at the final  $t = 2.0s$  timestep.  
 352

353 The performance of the baseline models highlights the advantages of our unified approach. T2P re-  
 354 sorts to an overly conservative strategy when challenged with this tricky, high-uncertainty scenario.  
 355 Its predictions become increasingly static over time, collapsing towards a mean pose with very lit-  
 356 tle movement to avoid large errors. In contrast, EMPMP attempts to generate dynamic motion but  
 357 struggles with physical plausibility. Its predictions exhibit noticeable artifacts, such as the unnat-  
 358ural arm posture of the person in green and the awkward leg movements of the person in blue. These  
 359 qualitative results underscore that our model not only achieves superior quantitative accuracy but  
 360 also produces motions that are significantly more realistic and coherent than competing methods.  
 361

Table 2: Throughput mean  $\pm$  std calculated over 10 runs on MOCAP-UMPM data. All models are run on a NVIDIA RTX A6000 GPU with batch size 64. Higher values are better ( $\uparrow$ ). An asterisk(\*) denotes models we recomputed for this experiment.

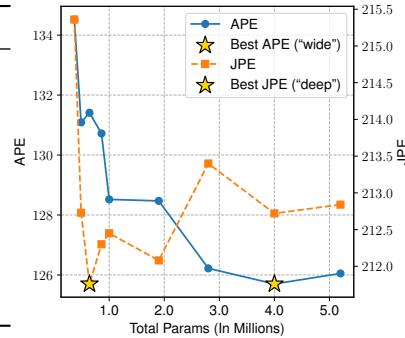
Model	Training Throughput (samples/sec) $\uparrow$	Test Throughput (samples/sec) $\uparrow$
T2P*	187 $\pm$ 22	401 $\pm$ 64
EMPMP*	812 $\pm$ 58	2041 $\pm$ 129
Ours (wide)	862 $\pm$ 43	2251 $\pm$ 140
Ours (deep)	<b>928 <math>\pm</math> 45</b>	<b>3673 <math>\pm</math> 161</b>



377 Figure 2: Visualization of predictions on a MOCAP-UMPM scene. Model predictions are in color,  
 378 and ground truth future poses are black dashes. The last-known input positions are colored dashes.  
 379

378 Table 3: Comparison of our model’s performance with different hyperparameter configurations.  
379

380 <b>Depth</b>	381 <b>Embed dim</b>	382 <b>Total Params</b>	383 <b>APE</b>	384 <b>JPE</b>
385 8	386 192	387 5.2M	388 126.05	389 212.84
390 6	391 192	392 4.0M	393 <b>125.70</b>	394 212.72
395 4	396 192	397 2.8M	398 126.22	399 213.40
400 12	401 96	402 1.9M	403 128.47	404 212.08
405 6	406 96	407 1.0M	408 128.52	409 212.45
410 12	411 64	412 860K	413 130.72	414 212.30
415 16	416 48	417 642K	418 131.41	419 <b>211.76</b>
420 12	421 48	422 490K	423 131.09	424 212.73
425 16	426 36	427 367K	428 134.52	429 215.36

391 

### 3.6 ABLATION STUDIES

393 In this section, we conduct a series of ablation studies to investigate the impact of our model’s key  
394 components and hyperparameters. We perform these experiments on the MOCAP-UMPM dataset  
395 for the joint pose and trajectory prediction task to analyze the effectiveness of our multi-modal  
396 prediction head and the trade-offs in our transformer architecture.

397 

#### 3.6.1 CHOICE OF TRANSFORMER HYPERPARAMETERS

399 Our model’s major computations are performed using a simple transformer decoder. We analyze the  
400 trade-offs between its depth (number of layers,  $L$ ) and width (embedding dimension,  $d_{\text{model}}$ ). We  
401 experimented with various configurations, keeping the overall parameter count relatively low, to find  
402 effective deep net architecture designs. The results are summarized in Table 3.

403 The analysis reveals a clear relationship between depth, width, and predictive accuracy. Our “wide”  
404 configuration ( $L = 6, d_{\text{model}} = 192$ ) achieves the best APE, suggesting that a more expansive  
405 embedding space is beneficial for capturing fine-grained pose details. Decreasing the depth to  $L = 4$   
406 or increasing it to  $L = 8$  with the same width leads to a decline in performance, indicating a sweet  
407 spot for this configuration.

408 Conversely, our “deep” model ( $L = 16, d_{\text{model}} = 48$ ) obtains the lowest JPE, demonstrating that a  
409 deeper stack of attention layers is more effective at modeling complex, long-range spatio-temporal  
410 dependencies for global trajectory prediction, even with a constrained embedding dimension. As  
411 expected, performance degrades significantly with shallower or narrower architectures. These  
412 results validate our choice of the “wide” and “deep” models, as they represent two distinct and highly  
413 effective points in the architecture design space, tailored for different aspects of motion prediction.

415 

#### 3.6.2 EFFECT OF MULTI-MODAL PREDICTION

416 While multi-modal prediction is standard in trajectory forecasting, state-of-the-art methods for joint  
417 pose and trajectory prediction, such as EMPMP, have often favored a deterministic approach, pre-  
418 dicting a single future outcome. However, human motion is inherently stochastic, and a single  
419 prediction can fail to capture the full range of plausible futures. We therefore conduct an ablation  
420 to quantify the advantage of our multi-modal prediction head explicitly. We compare our model’s  
421 performance when generating multiple proposals ( $K = 6$ ) against a deterministic setting ( $K = 1$ ),  
422 mirroring the setup of prior work (Jeong et al., 2024).

423 The results in Table 4 clearly demon-  
424 strate the limitations of a deterministic ap-  
425 proach. Even in a deterministic setting, our  
426 “wide” model is already competitive with  
427 EMPMP. However, by embracing multi-  
428 modality, our model achieves a dramatic  
429 performance gain. The APE improves by  
430 13.8% and the JPE by a substantial 24.2%.  
431 This highlights that our model doesn’t just  
432 produce a better single guess; it effectively

433 Table 4: Model performance with 2 different modes  
434 on MOCAP-UMPM data. Lower values are better (↓).

Model	$K = 1$		$K = 6$	
	APE	JPE	APE	JPE
T2P	154.4	366.4	151.7	262.7
EMPMP	147.2	283.1	146.5	250.4
Ours (wide)	<b>145.84</b>	<b>280.8</b>	<b>125.70</b>	212.72
Ours (deep)	149.35	286.96	131.41	<b>211.76</b>

captures a distribution of high-quality future motions. Interestingly, prior works do not benefit from multiple modes to the same degree. For instance, EMPMP’s APE barely improves, suggesting its architecture may struggle to generate genuinely distinct futures. While a full analysis of why the baselines are less suited to multi-modal prediction is beyond the scope of this paper, it suggests that our unified architecture is particularly effective at leveraging the “winner-takes-all” loss to produce a diverse and plausible set of outcomes—a crucial capability that deterministic models lack by design.

## 4 RELATED WORK

Human motion requires a holistic assessment, as local body articulation (pose) and global displacement (trajectory) are deeply intertwined. Our research community, however, has largely tackled motion prediction by decomposing this process into specialized sub-problems: pose, trajectory, and multi-person motion prediction. This specialization has driven progress on narrow benchmarks but created a dichotomy: specialized models fail to generalize, while the few holistic models struggle to compete on established task-specific leaderboards. This “benchmark effect” has incentivized an escalation in architectural complexity, with increasingly elaborate models gaining an edge.

In this context, the transformer has emerged as a powerful tool for sequence modeling. However, its application to human motion has often followed the increasing complexity trend, where it merely serves as a backbone for other domain-specific modules. This paper challenges that approach. We posit that the transformer’s true power lies not in its ability to support additional complex components, but in its inherent capacity to address the problem in a simple, direct, and unified manner.

### 4.1 HUMAN POSE PREDICTION

The task of human pose prediction involves forecasting a future sequence of 3D skeletal joint locations relative to the root joint based on an observed history of poses (Hosseininejad et al., 2025, see Appendix C). To address the stochastic nature of human behavior, the field has shifted from deterministic models (Medjaouri & Desai, 2022; Ma et al., 2022) to complex generative frameworks, particularly diffusion models. This pursuit of generative fidelity has fueled a cycle of escalating complexity, with methods like BeLFusion (Barquero et al., 2023) introducing a “behavioral latent space” and CoMusion (Sun & Chowdhary, 2024) employing a hybrid Transformer-GCN architecture that operates in the Discrete Cosine Transform (DCT) (Mao et al., 2021) space to model skeletal kinematics explicitly. Recent methods like SkeletonDiff (Curreli et al., 2025) and SLD (Xu et al., 2024) focus on skeleton-aware generation or long-sequence efficiency, while non-diffusion approaches like Motionmap (Hosseininejad et al., 2025) introduce novelties such as multi-stage heatmap pipelines.

### 4.2 HUMAN TRAJECTORY PREDICTION

Trajectory forecasting aims to predict the future path of an agent’s root joint, a task complicated by latent intent, social interactions, and environmental constraints. Recent state-of-the-art approaches have often relied on massive external knowledge sources or engineered, multi-stage pipelines. A prominent trend involves leveraging large foundation models; TrajCLIP (Yao et al., 2024), for example, incorporates knowledge from vision-language models (VLMs) to provide contextual cues, effectively outsourcing the learning problem. Another approach involves building complex frameworks for generality, such as Singular Trajectory (Bae et al., 2024), whose “universal” status is the result of an engineered pipeline involving Singular Value Decomposition and a diffusion-based refiner, or NMRF (Fang et al., 2025), which uses sophisticated modules like continuous, field-based scene representations.

### 4.3 COMBINED POSE AND TRAJECTORY PREDICTION

The simultaneous prediction of pose and trajectory is where the limitations of fragmented architectures become most apparent, as this task requires modeling the critical coupling between local articulation and global movement. Early work such as Tripod (Adeli et al., 2021) and work by Zaier et al. (2023) established the importance of forecasting these dynamics jointly, typically employing graph-based or multi-branch architectures to capture the dependencies. More recent approaches have explored pre-training strategies, such as Multi-transmition (Gao et al., 2024), to learn general-

486 **izable motion representations. Despite this progress, prior**Prior work has typically imposed strong  
 487 architectural priors on how pose and trajectory information should interact. T2P (Jeong et al., 2024)  
 488 employs a sequential, “coarse-to-fine” strategy, first predicting the global trajectory and then con-  
 489 ditioning the pose prediction on that result. This design imposes a one-way causal assumption that  
 490 trajectory dictates pose and is susceptible to error propagation. An alternative, seen in EMPMP  
 491 (Zheng et al., 2025), uses parallel branches to process local and global information separately be-  
 492 fore fusion. This avoids direct error propagation but imposes a prior: that local and global features  
 493 are separable concerns. This rigid separation may preclude the model from learning more complex,  
 494 deeply intertwined representations where local and global dynamics are jointly encoded from the  
 495 outset. Consequently, although EMPMP was explicitly designed to be “lightweight”, its architec-  
 496 ture is built from individually light but intricately integrated components and struggles to leverage  
 497 hardware parallelism effectively.  
 498

## 5 CONCLUSION

501 This paper introduces SimpliHuMoN, a simple and unified transformer-based model that addresses  
 502 the prevailing trends of fragmentation and escalating complexity in human motion prediction. We  
 503 challenge the field by demonstrating how a single, end-to-end framework effectively learns the dy-  
 504 namics of human movement across various tasks. Extensive experiments across a wide range of  
 505 standard benchmarks validate this approach, showing that our model achieves state-of-the-art accu-  
 506 racy while also proving more computationally efficient than prior methods. Ultimately, this work  
 507 serves as evidence that architectural simplicity, when thoughtfully applied, can outperform engi-  
 508 neered complexity, suggesting that the path forward in motion prediction lies not in adding more  
 509 intricate components but in refining simple and truly generalizable foundations.  
 510

## 6 REPRODUCIBILITY STATEMENT

512 We are committed to ensuring the reproducibility of our research. Our model’s architecture, loss  
 513 function, training procedure, and key hyperparameters are described in detail in Section 2 of the  
 514 main paper, with further analysis in our ablation studies (Section 3.6). For data handling, Ap-  
 515 pendix A provides a complete description of all datasets and the exact preprocessing steps, which  
 516 follow established protocols from prior work. The precise mathematical definitions for all evaluation  
 517 metrics are detailed in Appendix B. Finally, our supplementary website, referenced in Appendix C,  
 518 offers additional qualitative results. Collectively, these resources provide a comprehensive guide for  
 519 reproducing our experimental findings. We will also release all code.  
 520

## 521 REFERENCES

523 V. Adeli, M. Ehsanpour, I. D. Reid, J. C. Niebles, S. Savarese, E. Adeli, and H. Rezatofighi. Tripod:  
 524 Human trajectory and pose dynamics forecasting in the wild. In *Proc. of ICCV*, 2021.

525 I. Bae, J.H. Park, and H.G. Jeon. Learning pedestrian group representations for multi-modal trajec-  
 526 tory prediction. In *Proc. of ECCV*, 2022.

528 I. Bae, Y.J. Park, and H.G. Jeon. Singulartrajectory: Universal trajectory predictor using diffusion  
 529 model. In *Proc. of CVPR*, 2024.

530 G. Barquero, S. Escalera, and C. Palmero. Belfusion: Latent diffusion for behavior-driven human  
 531 motion prediction. In *Proc. of ICCV*, 2023.

533 N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and Zagoruyko S. End-to-end object  
 534 detection with transformers. In *Proc. of ICCV*, 2020.

535 A. Clark, A. W. Pillay, and D. Moodley. A system for pose analysis and selection in virtual reality  
 536 environments. In *Conference of the South African Institute of Computer Scientists and Infor-  
 537 mation Technologists*. ACM, 2020.

539 CMU Graphics Lab. CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu/>, 2003.

540 C. Curreli, D. Muhle, A. Saroha, Z. Ye, R. Marin, and D. Cremers. Nonisotropic Gaussian Diffusion  
 541 for Realistic 3D Human Motion Prediction. In *Proc. of CVPR*, 2025.

542

543 L. Dang, Y. Nie, C. Long, Q. Zhang, and G. Li. Diverse human motion prediction via gumbel-  
 544 softmax sampling from an auxiliary space. In *Proc. of ACM MM*, 2022.

545

546 Z. Fang, D. Hsu, and G. H. Lee. Neuralized markov random field for interaction-aware stochastic  
 547 human trajectory prediction. In *Proc. of ICLR*, 2025.

548

549 Q. Fu, X. Zhang, J. Xu, and H. Zhang. Capture of 3d human motion pose in virtual reality based on  
 550 video recognition. *Complexity*, 2020.

551

552 Y. Gao, P. Luan, and A. Alahi. Multi-transmotion: Pre-trained model for human motion prediction.  
 553 In *Proc. of CoRL*, 2024.

554

555 T. Gu, G. Chen, J. Li, C. Lin, Y. Rao, J. Zhou, and J. Lu. Stochastic trajectory prediction via motion  
 556 indeterminacy diffusion. In *Proc. of CVPR*, 2022.

557

558 R. Hosseininejad, M. Shukla, S. Saadatnejad, M. Salzmann, and A. Alahi. MotionMap: Represent-  
 559 ing Multimodality in Human Pose Forecasting. In *Proc. of CVPR*, 2025.

560

561 C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and  
 562 predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern  
 563 Analysis and Machine Intelligence*, 36(7), 2013.

564

565 J. Jeong, D. Park, and K.J. Yoon. Multi-agent long-term 3d human pose forecasting via interaction-  
 566 aware trajectory conditioning. In *Proc. of CVPR*, 2024.

567

568 T. Jiang, J. Billingham, S. Müksch, J. Zarate, N. Evans, M. Oswald, M. Pollefeys, O. Hilliges,  
 569 M. Kaufmann, and J. Song. Worldpose: A world cup dataset for global 3d human pose estimation.  
 570 In *Proc. of ECCV*, 2024.

571

572 A. Lerner, Y. Chrysanthou, and D. Lischinski. Crowds by example. In *Proc. of CVPR*, 2007.

573

574 B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning.  
 575 In *Proc. of EMNLP*, 2021.

576

577 Y. Li, C. Chang, C. Cheng, and Y. Huang. Baseball swing pose estimation using openpose. In *Proc.  
 578 of ICRA*, 2021.

579

580 T. Ma, Y. Nie, C. Long, Q. Zhang, and G. Li. Progressively generating better initial guesses towards  
 581 next stages for high-quality human motion prediction. In *Proc. of CVPR*, 2022.

582

583 N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black. Amass: Archive of motion  
 584 capture as surface shapes. In *Proc. of ICCV*, 2019.

585

586 W. Mao, M. Liu, and M. Salzmann. Generating smooth pose sequences for diverse human motion  
 587 prediction. In *Proc. of ICCV*, 2021.

588

589 O. Medjaouri and K. Desai. HR-STAN: High-Resolution Spatio-Temporal Attention Network for  
 590 3D Human Motion Prediction. In *Proc. of CVPR*, 2022.

591

592 B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli. A survey of motion planning and control  
 593 techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1), 2016.

594

595 S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You'll Never Walk Alone: Modeling Social  
 596 Behavior for Multi-target Tracking. In *Proc. of ICCV*, 2009.

597

598 H. Ro, Y. J. Park, J.-H. Byun, and T.-D. Han. Display methods of projection augmented reality  
 599 based on deep learning pose estimation. In *ACM SIGGRAPH Posters*, 2019.

600

601 A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory  
 602 understanding in crowded scenes. In *Proc. of ECCV*, 2016.

594 T. Salzmann, H.-T. L. Chiang, M. Ryll, D. Sadigh, C. Parada, and A. Bewley. Robots that can see:  
 595 Leveraging human pose for trajectory prediction. *IEEE Robotics and Automation Letters*, 2023.  
 596

597 L. Shi, L. Wang, S. Zhou, and G. Hua. Trajectory unified transformer for pedestrian trajectory  
 598 prediction. In *Proc. of ICCV*, 2023.

599 J. Sun and G. Chowdhary. Comusion: Towards consistent stochastic human motion prediction via  
 600 motion diffusion. In *Proc. of ECCV*, 2024.

601

602 N. P. van der Aa, X. Luo, G. J. Giezeman, R. T. Tan, and R. C. Veltkamp. Umpm benchmark: A  
 603 multi-person dataset with synchronized video and motion capture data for evaluation of articulated  
 604 human motion and interaction. In *Proc. of ICCV Workshops*, 2011.

605 T. von Marcard, R. Henschel, M. J. Black, B. Rosenhahn, and G. Pons-Moll. Recovering accurate  
 606 3d human pose in the wild using imus and a moving camera. In *Proc. of ECCV*, 2018.

607

608 G. Xu, J. Tao, W. Li, and L. Duan. Learning semantic latent directions for accurate and controllable  
 609 human motion prediction. In *Proc. of ECCV*, 2024.

610

611 P. Yao, Y. Zhu, H. Bi, T. Mao, and Z. Wang. Trajclip: Pedestrian trajectory prediction method using  
 612 contrastive learning and idempotent networks. In *Advances in Neural Information Processing  
 Systems*, 2024.

613

614 M. Zaier, H. Wannous, H. Drira, and J. Boonaert. A dual perspective of human motion analysis - 3d  
 615 pose estimation and 2d trajectory prediction. In *Proc. ICCV Workshops*, 2023.

616

617 J. Zheng, X. Shi, A. Gorban, J. Mao, Y. Song, C. R. Qi, T. Liu, V. Chari, A. Cornman, Y. Zhou,  
 618 C. Li, and D. Anguelov. Multi-modal 3d human pose estimation with 2d weak supervision in  
 619 autonomous driving. In *Proc. of CVPR*, 2022.

620

621 Y. Zheng, R. Yu, and J. Sun. Efficient multi-person motion prediction by lightweight spatial and  
 622 temporal interactions. In *Proc. of ICCV*, 2025.

623

624 J. Zou. Simplified neural architecture for efficient human motion prediction in human-robot inter-  
 625 action. *Neurocomputing*, 588, 2024.

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 APPENDIX: SIMPLIHUMON: SIMPLIFYING HUMAN MOTION PREDICTION  
649650 This appendix is structured as follows: In Sec. A we provide additional dataset and metric details. In  
651 Sec. B we detail additional experimental results. In Sec. C we highlight the website which is part of  
652 the provided appendix. In Sec. D we discuss joint training results. In Sec. E we provide information  
653 about our LLM usage.  
654655 A ADDITIONAL DATASET AND METRIC DETAILS  
656657 A.1 SOURCES AND PROCESSING OF DATA  
658659 All experiments are conducted on publicly available, open-source datasets. To ensure a fair and  
660 direct comparison with prior work, we strictly adhere to the established data processing and eval-  
661 uation protocols from recent top-performing methods for each prediction task. This standardization  
662 ensures that the performance improvements reported in this paper are attributable to our model’s  
663 architecture rather than differences in data handling. The specific protocols are as follows: For pose  
664 prediction on the Human3.6M and AMASS datasets, we follow the data processing methodology,  
665 sequence lengths, and evaluation splits established by BeLFusion (Barquero et al., 2023). For trajec-  
666 tory prediction on the ETH-UCY and SDD, our data handling and evaluation procedures align with  
667 the protocol set forth by NMRF (Fang et al., 2025). For combined pose and trajectory Prediction  
668 on the MOCAP-UMPM and 3DPW datasets, we adopt the data preparation and processing pipeline  
669 outlined by T2P (Jeong et al., 2024).  
670

## 671 A.2 METRIC FORMULAE

672 Given the predicted motion proposal  $X_{\text{fut}}^k = \{x_{t,m}^k\} \in \mathbb{R}^{F \times M \times 3}$  for  $k \in \{1, 2, \dots, K\}$  across  $F$   
673 time frames with  $M$  joints per person, along with the corresponding ground truth  $X_{\text{fut}}^{\text{gt}} = \{x_{t,m}^{\text{gt}}\}$ , the  
674 following metrics are used for evaluation. For multi-modal predictions, we follow common practice  
675 and report the minimum error among all  $K$  generated proposals for each metric (e.g., minADE,  
676 minFDE). Consistent with prior work, for datasets containing multiple people, the final reported  
677 error is the average of the metric computed for all individuals. All metrics in the main paper are  
678 reported for the final output timestep,  $t = F$ .  
679680 **APE.** Aligned mean per joint Position Error (APE) is used as a metric to evaluate the forecasted  
681 local motion. Euclidean distance of each joint relative to the root (hip) joint is averaged over all  
682 joints for a given timestep,  $t$ :  
683

684 
$$\text{APE}_t(X_{\text{fut}}^{\text{gt}}, X_{\text{fut}}^k) = \frac{1}{M} \sum_{m=1}^M \|(x_{t,m}^{\text{gt}} - x_{t,\text{hip}}^{\text{gt}}) - (x_{t,m}^k - x_{t,\text{hip}}^k)\|_2. \quad (2)$$
  
685

686 **JPE.** Joint Precision Error (JPE) evaluates both global and local predictions by the average Eu-  
687 clidean distance of all joints for a given timestep,  $t$ :  
688

689 
$$\text{JPE}_t(X_{\text{fut}}^{\text{gt}}, X_{\text{fut}}^k) = \frac{1}{M} \sum_{m=1}^M \|x_{t,m}^{\text{gt}} - x_{t,m}^k\|_2. \quad (3)$$
  
690

692 **ADE.** Average Displacement Error (ADE) measures the Euclidean distance between the ground  
693 truth and predicted sequences, averaged over all joints and all future time frames:  
694

695 
$$\text{ADE}(X_{\text{fut}}^{\text{gt}}, X_{\text{fut}}^k) = \frac{1}{F \times M} \sum_{t=1}^F \sum_{m=1}^M \|x_{t,m}^{\text{gt}} - x_{t,m}^k\|_2. \quad (4)$$
  
696

698 **FDE.** Final Displacement Error (FDE) measures the Euclidean distance between the ground truth  
699 and the prediction, averaged over all joints for a given timestep,  $t$ :  
700

701 
$$\text{FDE}_t(X_{\text{fut}}^{\text{gt}}, X_{\text{fut}}^k) = \frac{1}{M} \sum_{m=1}^M \|x_{t,m}^{\text{gt}} - x_{t,m}^k\|_2. \quad (5)$$

702 Table 5: Trajectory prediction performance (ADE/FDE) on ETH-UCY. Lower values are better, with  
 703 the best results shown in **bold**. A dagger ( $\dagger$ ) marks models adapted for the specific task, while a ( $\wedge$ )  
 704 notes models that use external training data.

706 <b>Model</b>	707 <b>ETH</b>	708 <b>HOTEL</b>	709 <b>UNIV</b>	710 <b>ZARA1</b>	711 <b>ZARA2</b>	712 <b>AVG</b>
MID	0.39/0.66	0.13/0.22	0.22/0.45	0.17/0.30	0.13/0.27	0.21/0.38
GP-Graph	0.43/0.63	0.18/0.30	0.24/0.42	0.17/0.31	0.15/0.29	0.23/0.39
TUTR	0.40/0.61	0.11/0.18	0.23/0.42	0.18/0.34	0.13/0.25	0.21/0.36
SingularTrajectory	0.35/0.42	0.13/0.19	0.25/0.44	0.19/0.32	0.15/0.25	0.22/0.34
TrajCLIP $\wedge$	0.36/0.57	<b>0.10/0.17</b>	<b>0.19/0.41</b>	<b>0.16/0.28</b>	<b>0.11/0.20</b>	<b>0.18/0.33</b>
NMRF	<b>0.26/0.37</b>	0.11/0.17	0.28/0.49	0.17/0.30	0.14/0.25	0.19/ <b>0.32</b>
T2P $\dagger$	0.29/0.55	0.15/0.27	0.25/0.53	<b>0.16/0.33</b>	0.12/0.26	0.19/0.39
EMPMP $\dagger$	0.99/0.98	0.70/0.87	0.69/0.89	0.43/0.50	0.32/0.35	0.63/0.72
Ours (wide)	0.28/0.44	0.13/0.24	0.24/0.44	<b>0.16/0.29</b>	<b>0.11/0.21</b>	<b>0.18/0.32</b>
Ours (deep)	0.29/0.44	0.14/0.24	0.24/0.43	0.17/0.29	0.13/0.21	0.19/ <b>0.32</b>

## 718 B ADDITIONAL EXPERIMENTAL RESULTS

### 721 B.1 PER-DATASET SPLIT ON ETH-UCY

722 On the ETH-UCY datasets, our model demonstrates highly competitive performance against leading  
 723 methods, as detailed in Table 5. While models like TrajCLIP (Yao et al., 2024) and NMRF (Fang  
 724 et al., 2025) achieve the best results on some of the individual scenes, our “wide” configuration  
 725 achieves the best overall performance, tying for the best average ADE (0.18) and the best average  
 726 FDE (0.32).

727 This result is particularly noteworthy when considering the architectural differences between our  
 728 model and methods like TrajCLIP. TrajCLIP’s strong performance stems from its use of a large,  
 729 pre-trained VLM to provide rich semantic priors. Specifically, it uses natural language prompts  
 730 (*e.g.*, “a person walking”) to generate contextual embeddings from the VLM’s text encoder, which  
 731 are then fused with visual features to guide the trajectory prediction. This approach effectively  
 732 outsources a part of the learning problem to a massive external knowledge base. While powerful,  
 733 this creates a dependency on computationally heavy external models and assumes that general web-  
 734 scale knowledge is optimally suited for the fine-grained physics of trajectory prediction.

735 Our model, in contrast, is entirely self-contained, learning all necessary dynamics exclusively from  
 736 the provided motion data. The performance difference on the ETH scene, where our model signif-  
 737 icantly outperforms TrajCLIP, suggests a key advantage of this self-sufficient approach. The ETH  
 738 dataset represents a scenario where the visual-semantic cues that TrajCLIP relies on are less infor-  
 739 mative and reliable than in other scenes. In such cases, our model’s ability to learn robustly from  
 740 the motion dynamics alone allows it to generalize more effectively, leading to a more consistent  
 741 performance profile across all five datasets. This consistency is what enables our model to achieve  
 742 better average performance without relying on external priors, challenging the notion that they are a  
 743 prerequisite for top-tier trajectory forecasting.

744 Furthermore, a key architectural difference is TrajCLIP’s explicit modeling of social and environ-  
 745 mental interactions through two dedicated modules. They are designed to capture the dynamics  
 746 between different agents and integrate visual context from the environment to make predictions  
 747 physically consistent with the static scene. In contrast, our current model processes each agent  
 748 independently and contains no such explicit interaction mechanisms. The fact that our simpler,  
 749 non-interactive approach still achieves state-of-the-art average performance highlights the remarkable  
 750 strength and efficiency of its core motion representation. This also points to a promising avenue for  
 751 future work: integrating a lightweight interaction mechanism into our powerful architecture could  
 752 potentially push performance even further.

756  
757  
758  
759 Table 6: Comparison of APE/JPE metrics across models and datasets. Lower values are better ( $\downarrow$ ),  
760 with the best results shown in **bold**. An asterisk (\*) denotes models we recomputed for this setup.  
761  
762  
763  
764  
765

766	767	768	769	In/Out Length (s)	MOCAP-UMPM					3DPW			
					770 0.4s	0.8s	1.2s	1.6s	2.0s	771 0.4s	0.8s	1.2s	1.6s
772 <b>APE</b>	T2P*	773	774 71.7	775 107.8	776 120.4	777 137.1	778 151.7	779 98.2	780 114.6	781 135.3	782 150.0	783	784
	EMPMP*	785	786 60.1	787 96.0	788 116.9	789 131.6	790 146.5	791 96.3	792 111.9	793 134.4	794 150.6	795	796
	Ours (wide)	797	798 <b>57.3</b>	799 <b>87.7</b>	800 <b>104.5</b>	801 <b>115.3</b>	802 <b>125.7</b>	803 92.8	804 <b>107.1</b>	805 <b>130.0</b>	806 <b>142.9</b>	807	808
	Ours (deep)	809	810 62.3	811 89.5	812 107.3	813 119.0	814 128.5	815 <b>93.2</b>	816 108.4	817 131.5	818 148.9	819	820
772 <b>JPE</b>	T2P*	773	774 70.2	775 139.2	776 160.1	777 226.4	778 262.7	779 107.7	780 142.6	781 181.0	782 236.2	783	784
	EMPMP*	785	786 68.0	787 123.9	788 170.3	789 219.1	790 250.4	791 103.6	792 140.2	793 179.8	794 235.4	795	796
	Ours (wide)	797	798 <b>64.6</b>	799 <b>108.6</b>	800 <b>143.9</b>	801 <b>177.7</b>	802 <b>212.7</b>	803 <b>99.4</b>	804 <b>137.3</b>	805 <b>172.1</b>	806 <b>231.0</b>	807	808
	Ours (deep)	809	810 68.9	811 109.9	812 145.3	813 <b>177.2</b>	814 <b>210.3</b>	815 100.1	816 138.2	817 <b>171.6</b>	818 231.5	819	820

## 771 B.2 DETAILED METRICS ACROSS KEY FRAMES

772 To scrutinize performance over the forecast horizon, Table 6 presents a time-step-level analysis on  
773 the MOCAP-UMPM and 3DPW datasets. The results reveal not only the consistent superiority of  
774 our models over T2P and EMPMP at every interval but also a crucial architectural trade-off.  
775

776 Our “wide” model establishes a new standard for local pose accuracy (APE), excelling at capturing  
777 fine-grained kinematics, particularly in the short term. Conversely, our “deep” model demonstrates  
778 its strength in long-range forecasting, achieving the best overall world-coordinate accuracy (JPE) at  
779 the final timesteps. This divergence highlights a key finding: architectural depth appears more criti-  
780 cal for maintaining global trajectory coherence, while width is more effective for local pose detail.  
781 Most notably, the performance gap between our models and the baselines widens as the prediction  
782 horizon increases. This demonstrates our architecture’s superior robustness against the error accu-  
783 mulation that typically plagues sequential prediction tasks. This detailed analysis confirms that our  
784 simple, unified framework is not just more accurate overall but is also more effective at handling  
785 the challenges of long-term motion forecasting compared to competing multi-stage or specialized  
786 approaches.

## 787 C WEBSITE

788 We provide a website with additional visualizations demonstrating our method’s performance, which  
789 can be accessed using the provided HTML file.  
790

791 We observe that the generated motions exhibit high physical plausibility, with no unrealistic arti-  
792 facts such as foot sliding. Body poses are consistently realistic, respecting natural body constraints  
793 and capturing fine-grained details without grouping different joints into unnatural, blocky move-  
794 ments. Furthermore, our model adeptly handles both independent and coupled motion dynamics; it  
795 accurately predicts localized movements (*e.g.*, arm gestures without a change in trajectory) and com-  
796 plex actions where limb articulation and global trajectory are deeply intertwined. Our model excels  
797 in multi-person scenes by processing agents independently. This avoids a key limitation of rigid,  
798 graph-based interaction models (GNNs), which can corrupt individual forecasts by forcing informa-  
799 tion aggregation from non-interacting neighbors. This finding does not diminish the importance of  
800 interaction modeling but rather clarifies the need to learn it dynamically.  
801

## 802 D JOINT TRAINING

803 To test the full generalization capability of our architecture, we train a single, universal model jointly  
804 on all datasets across all tasks (pose, trajectory, and combined prediction). This experiment aims  
805 to create a single set of weights that can perform any of the specialized tasks without retraining.  
806 Handling the significant diversity in data formats, skeleton structures, and sequence lengths requires  
807 a carefully designed methodology, which we detail below.  
808

810  
 811 Table 7: Mapping from dataset-specific skeletons to our 22-joint canonical representation. AMASS  
 812 serves as the canonical skeleton itself. Dashes (–) indicate that a direct mapping for that specific  
 813 canonical joint is unavailable in the source dataset.

#	AMASS	Human3.6M	MOCAP-UMPM	3DPW
1	Pelvis	–	Hips	Pelvis
2	L_Hip	LeftUpLeg	LHip	LHip
3	R_Hip	RightUpLeg	RHip	RHip
4	Spine1	Spine	Spine	–
5	L_Knee	LeftLeg	LKnee	LKnee
6	R_Knee	RightLeg	RKnee	RKnee
7	Spine2	–	–	–
8	L_Ankle	LeftFoot	LAngle	–
9	R_Ankle	RightFoot	RAngle	–
10	Spine3	–	–	–
11	L_Foot	–	–	LFoot
12	R_Foot	–	–	RFoot
13	Neck	Neck	Neck	–
14	L_Collar	–	–	–
15	R_Collar	–	–	–
16	Head	Head / Head-top	Head	–
17	L_Shoulder	LeftArm	LShoulder	LShoulder
18	R_Shoulder	RightArm	RShoulder	RShoulder
19	L_Elbow	LeftForeArm	LElbow	LElbow
20	R_Elbow	RightForeArm	RElbow	RElbow
21	L_Wrist	LeftHand	LWrist	LWrist
22	R_Wrist	RightHand	–	RWrist

### D.1 METHODOLOGY

819  
 820 **Data Unification and Canonical Skeleton.** A primary challenge is the heterogeneity of the  
 821 datasets. To create a consistent input format, all data is preprocessed into a normalized tensor of  
 822 shape  $T \times M \times 3$  (sequence length  $\times$  joints  $\times$  coordinates). We pad the data with a zero Z-dimension  
 823 for 2D trajectory datasets (ETH-UCY, SDD) to create a consistent 3D representation.

824 To address the varying skeleton definitions, we establish a 22-joint canonical skeleton, using the  
 825 AMASS dataset as our standard. All other datasets are mapped to this representation, as shown in  
 826 Table 7. This mapping allows us to use a fixed set of learnable joint embeddings, ensuring that input  
 827 data for a given semantic body part (e.g., the ‘Left Knee’) is always processed by its corresponding  
 828 embedding, regardless of the source dataset. For trajectory-only datasets, the single trajectory point  
 829 is mapped to the ‘Pelvis’ joint embedding.

830 **Dataset-Balanced Batching.** We employ a dataset-balanced batching strategy to prevent the  
 831 model from overfitting to larger datasets (e.g., AMASS). Each training batch contains samples drawn  
 832 from only a single dataset. We iterate through an equal number of batches from every dataset during  
 833 each epoch, ensuring the model is exposed to a balanced distribution of tasks and data sources  
 834 during training.

835 **Task-Specific Processing.** We use a task-type flag associated with each dataset to direct samples  
 836 through the appropriate processing pipelines. For instance, a ‘trajectory’ flag ensures that data only  
 837 passes through the trajectory-related input and output heads of the model, while a ‘joint’ flag ac-  
 838 tivates both pose and trajectory heads. This allows the shared transformer core to learn a general  
 839 motion representation while the specialized heads handle the task-specific details.

840 **Unified Model with Dynamic Slicing.** The model’s internal parameters are defined by the max-  
 841 imum sequence length,  $\max(T)$ , and maximum number of joints,  $\max(M)$ , across all datasets.  
 842 However, at runtime, a given sample’s input and output tensors are dynamically sliced to match the

864 Table 8: Comparison of performance on individual vs. joint training. Lower values are better (↓).  
865

866 Dataset 867 In/Out length (s) 868 Metric	869 Pose Prediction		870 Trajectory Prediction		871 Pose + Trajectory Prediction	
	872 Human3.6M 0.5/2.0 873 ADE↓/FDE↓	874 AMASS 0.5/2.0 875 ADE↓/FDE↓	876 ETH-UCY (Avg) 3.2/4.8 877 ADE↓/FDE↓	878 SDD 3.2/4.8 879 ADE↓/FDE↓	880 MOCAP-UMPM 1.0/2.0 881 APE↓/JPE↓	882 3DPW 0.8/1.6 883 APE↓/JPE↓
	Ours (wide, ind.)	0.42/0.59	0.31/0.45	0.18/0.32	6.70/7.63	125.70/212.72
Ours (deep, ind.)	0.44/0.57	0.35/0.47	0.19/0.32	6.26/7.61	131.41/211.76	148.91/231.48
Ours (wide, joint)	0.49/0.63	0.51/0.66	0.23/0.37	9.04/11.21	135.19/220.13	150.40/234.81
Ours (deep, joint)	0.55/0.70	0.62/0.78	0.25/0.39	10.66/12.14	138.20/223.49	151.46/235.05

874 specific  $T$  and  $M$  of its source dataset. This allows a single, fixed-size model to efficiently process  
875 variable-dimension inputs and outputs.

## 877 D.2 RESULTS

879 The results of our joint training experiment, presented in Table 8, demonstrate both the promise  
880 and the challenges of creating a single, universal motion prediction model. As expected, there is a  
881 performance trade-off when compared to the specialized models trained on individual datasets. The  
882 jointly trained models exhibit a degradation in accuracy across all tasks and datasets. However, the  
883 degree of this degradation varies, providing valuable insights into the model’s behavior.

884 The “wide” model consistently outperforms the “deep” model in the joint training setting. This  
885 is the inverse of our findings in some specialized tasks, and it suggests that the higher parameter  
886 count and wider embedding dimension of the “wide” model provide the necessary capacity to learn  
887 a shared representation across the seven diverse datasets. The “deep” model, with its constrained  
888 architecture, likely lacks the capacity to effectively generalize across such a heterogeneous data  
889 distribution, leading to a more significant performance drop. We also observe that the performance  
890 degradation is most pronounced on the AMASS dataset. This is likely a direct consequence of our  
891 dataset-balanced batching strategy. While this strategy prevents the model from overfitting to the  
892 largest datasets, it also means that the model is significantly under-exposed to the vast and diverse  
893 AMASS dataset, which is over 140 times larger than the smallest dataset (SDD). The model simply  
894 does not see enough of the AMASS data distribution to learn it as effectively as the specialized  
895 model.

896 Despite the performance trade-off, these results represent a successful proof of concept. The ability  
897 of a single, simple architecture to perform pose prediction, trajectory forecasting, and combined  
898 holistic prediction without any architectural changes is a powerful demonstration of its inherent  
899 generality. The fact that the model produces reasonable, albeit less accurate, predictions across all  
900 tasks indicates that it has learned a meaningful and transferable internal representation of human  
901 motion. This experiment validates the potential for developing true “foundation models for motion.”  
902 While our current approach shows a performance gap, it highlights a clear and promising research  
903 direction. Future work could focus on more sophisticated data-balancing techniques, curriculum  
904 learning strategies, or simply scaling the model’s capacity to bridge this gap. The ability to train  
905 a single model that understands the principles of human motion across myriad contexts remains a  
906 valuable and achievable goal for the field.

## 907 E LLM USAGE

909 While preparing this work, we used an LLM to assist with language editing and code generation  
910 for LaTeX tables and visualizations. The LLM’s contributions were limited to improving the clarity  
911 of the text and formatting results. The core research, experimental design, and all scientific claims  
912 remain our original work.

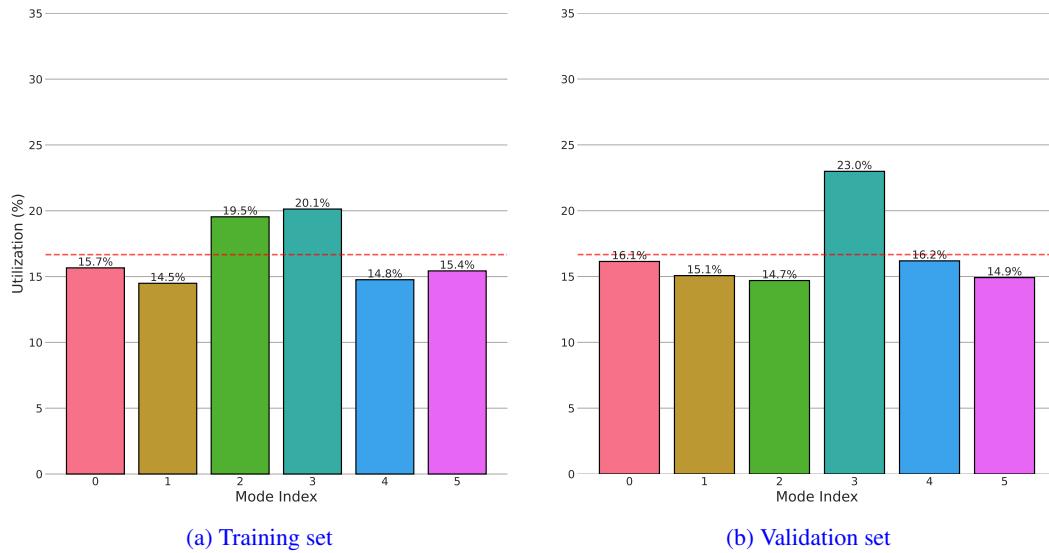


Figure 3: Distribution of winning mode indices (best-of-6) on pose + trajectory prediction task across the training and validation sets of MOCAP-UMPM. The dashed line (---) indicates equal distribution. Both distributions demonstrate balanced mode utilization without mode collapse.

## F EVALUATING PREDICTION DIVERSITY

A key component of our model is the multi-modal prediction head, which generates  $K$  distinct hypotheses to account for the uncertain nature of human motion. To validate its effectiveness, we analyze two potential concerns: mode collapse and the true diversity of the generated futures.

### F.1 MODE UTILIZATION ANALYSIS

To quantitatively verify that our model uses its full predictive capacity, we logged the index of the best (lowest error) hypothesis for every sample in the MOCAP-UMPM training and validation sets ( $K = 6$ ). Figure 3 plots the distribution of these winning indices. The results show that the model does not suffer from mode collapse. All six modes are actively utilized in both training and validation, with utilization rates clustering around the ideal uniform distribution (16.7%, shown as a dashed line). This confirms that the “winner-takes-all” loss, when applied to our architecture, successfully encourages the different proposals to specialize and cover distinct, plausible outcomes.

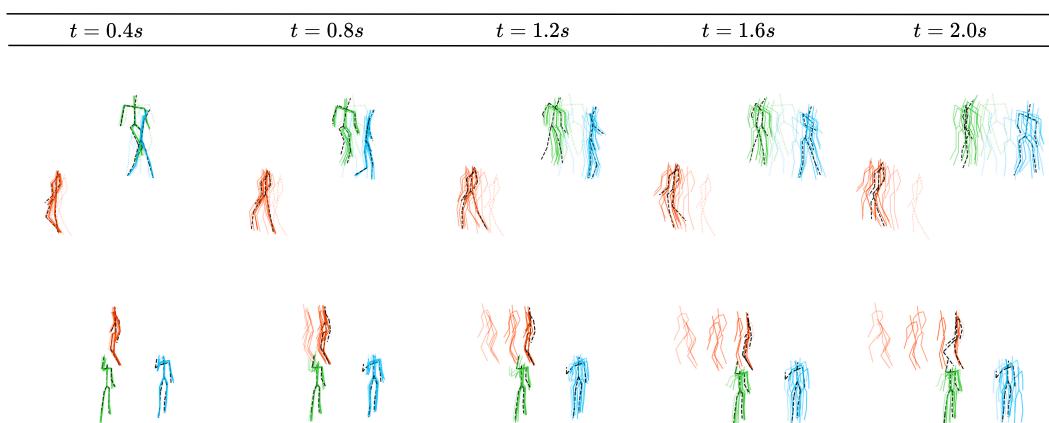


Figure 4: Visualization of motion proposals ( $K = 6$ ) of our (wide) model on MOCAP-UMPM data. All model predictions are in color. Ground truth future poses are black dashes, and the last-known input positions are colored dashes.

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
Table 9: Comparison of pose prediction diversity. Lower values are better ( $\downarrow$ ), with the best results shown in **bold**. An asterisk (\*) denotes models we recomputed for this setup.

Model	Type	Human3.6M		AMASS	
		MMADE $\downarrow$	MMFDE $\downarrow$	MMADE $\downarrow$	MMFDE $\downarrow$
DivSamp	Stochastic (Gumbel-Softmax)	0.542	0.671	0.623	0.728
BeLFusion	Stochastic (Latent Diffusion)	0.491	0.586	0.488	0.564
CoMusion	Stochastic (Motion Diffusion)	0.531	0.623	0.526	0.602
Motionmap	Stochastic (Multi-Stage Encoder-Decoder)	<b>0.466</b>	<b>0.532</b>	<b>0.450</b>	<b>0.514</b>
SkeletonDiff*	Stochastic (Gaussian Diffusion)	0.568	0.694	0.641	0.740
SLD*	Stochastic (State-Space Diffusion)	0.497	0.576	0.482	0.551
Ours (wide)	Deterministic (K-Proposal)	0.526	0.587	0.519	0.560
Ours (deep)	Deterministic (K-Proposal)	0.535	0.597	0.521	0.571

## F.2 QUALITATIVE DIVERSITY VISUALIZATION

Figure 4 provides a qualitative visualization of this diversity for a sample from the MOCAP-UMPM dataset. The figure overlays all  $K = 6$  proposals generated by our “wide” model. It clearly illustrates that the model is capturing distinct, high-level behaviors; for instance, the agent in red is predicted to either walk straight, stop, or turn, with each prediction representing a physically plausible and coherent motion.

## F.3 QUANTITATIVE DIVERSITY METRICS

To further assess the quality and diversity of our generated motion distribution, we report standard multimodal metrics, Minimum-over-K Average Displacement Error (MMADE) and Minimum-over-K Final Displacement Error (MMFDE), on the pose-only benchmarks. Table 9 compares our model’s performance against prominent stochastic and generative baselines. Our deterministic K-proposal approach achieves MMADE/MMFDE scores that are competitive with these generative methods, demonstrating that our  $K$  proposals capture a meaningful and diverse set of high-quality future motions.

## G ANALYSIS OF UNIFIED ARCHITECTURE

To validate our central claims, we present ablations addressing two critical questions. First, we provide quantitative proof that jointly modeling pose and trajectory is mutually beneficial. Second, we investigate why our simple, unified attention mechanism outperforms standard, more complex encoder-decoder designs. We also extend our ablation study from Section 3.6 to test the inclusion of some architectural components.

### G.1 BENEFIT OF JOINT MODELING

A core hypothesis of our work is that pose and trajectory are deeply intertwined and that modeling them jointly improves the prediction of both. We tested this hypothesis directly by training “pose-only” and “trajectory-only” variants of our model and comparing them to our full, joint model on MOCAP-UMPM.

Table 10 presents these results, providing strong quantitative evidence for our hypothesis. The data shows that pose prediction improves by  $\sim 11\text{-}12\%$  when trajectory is provided as an input, compared to a “pose-only” variant. Conversely, trajectory prediction improves by  $\sim 12\text{-}14\%$  when pose is provided as an input, compared to a “trajectory-only” variant. This confirms that our unified framework successfully learns the physical coupling between local articulation (pose) and global movement (trajectory), using information from one to refine its predictions for the other.

### G.2 UNIFIED SELF-ATTENTION VS. CROSS-ATTENTION

A key question is why our simple architecture works so well. We hypothesize it is due to our unified self-attention mechanism, where past context and future queries are concatenated as  $([\mathcal{C};$

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079  
Table 10: Comparison of the influence of combined pose and trajectory information vs. individual performance on MOCAP-UMPM data. Lower values are better ( $\downarrow$ ), with the best results shown in **bold**. A dagger ( $\dagger$ ) marks models adapted for the specific task.

Training Method Metric	Pose Prediction		Trajectory Prediction	
	Only Pose ADE $\downarrow$ /FDE $\downarrow$	Pose + Traj ADE $\downarrow$ /FDE $\downarrow$	Only Traj ADE $\downarrow$ /FDE $\downarrow$	Pose + Traj ADE $\downarrow$ /FDE $\downarrow$
T2P $\dagger$	-	0.61/0.95	-	0.20/0.29
EMPMP $\dagger$	-	0.53/0.65	-	0.14/0.20
Ours (wide)	0.46/0.58	<b>0.41/0.51</b>	0.09/0.18	<b>0.08/0.16</b>
Ours (deep)	0.46/0.59	0.42/0.51	0.09/0.18	<b>0.08/0.17</b>
Improvement (%)		<b>-11.4/-12.2</b>		<b>-13.7/-11.9</b>

Table 11: Ablation study on attention mechanisms measuring metrics on MOCAP-UMPM. Lower values are better ( $\downarrow$ ), with the best results shown in **bold**.

Attention Mechanism	Wide		Deep	
	APE $\downarrow$	JPE $\downarrow$	APE $\downarrow$	JPE $\downarrow$
Self-Attn over $[\mathcal{C}; \mathcal{Q}]$ (Ours)	<b>125.70</b>	<b>212.72</b>	<b>131.41</b>	<b>211.76</b>
Self-Attn ( $\mathcal{C}$ ) + Cross ( $\mathcal{Q}$ )	134.61	229.05	140.32	227.89

$\mathcal{Q}]$ ) and processed in a single attention block. This differs from standard encoder-decoders that use separate self-attention on the context and cross-attention for the queries. To test this, we built a new baseline that replaces our unified attention with a standard encoder-decoder design, keeping all other parameters and hyperparameters identical.

Table 11 shows the results. Our unified Self-Attn( $[\mathcal{C}; \mathcal{Q}]$ ) mechanism outperforms the standard cross-attention baseline, improving APE by 6.6% and JPE by 7.1% for the wide model. We believe this is because unified attention allows for a richer, bidirectional information flow at every step.

Figure 5 provides a visualization of the attention patterns in our first transformer layer. The heatmap shows attention in all four quadrants, representing the bidirectional interactions between past context  $[\mathcal{C}]$  and future query  $[\mathcal{Q}]$  tokens. Brighter colors, indicating stronger attention, are visible for queries attending to the past context (bottom-left quadrant) but also for queries attending to other queries (bottom-right). This explicitly shows the model is learning the complex, bidirectional relationships and not just the standard query-to-context flow, enabled by our unified attention mechanism.

This visual evidence supports our hypothesis that this richer, bidirectional attention is key to our model’s effectiveness.

### G.3 ARCHITECTURAL COMPONENT ABLATION

We investigate the contribution of our smaller architectural choices. Table 12 analyzes the impact of removing the Type Embedding ( $\epsilon$ ) and replacing RMSNorm with LayerNorm on the MOCAP-

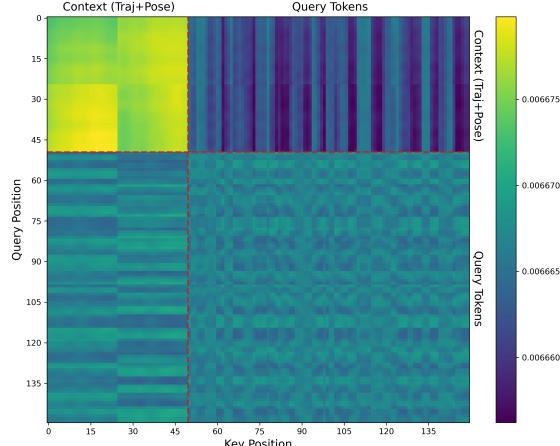
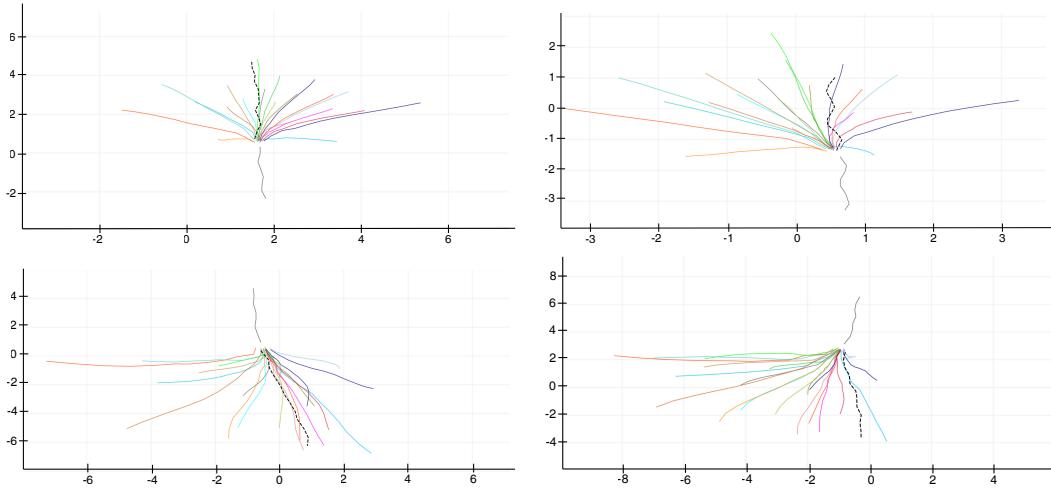


Figure 5: Attention patterns in the first transformer block at epoch 100. Brighter colors indicate stronger attention weights. Dashed lines (---) separate past context from future query tokens.

1080  
1081 Table 12: Extended ablation study of architecture choices on MOCAP-UMPM. We explicitly com-  
1082 pare RMSNorm vs. LayerNorm and the effect of Type Embeddings. Lower values are better ( $\downarrow$ ),  
1083 with the best results shown in **bold**.

Components			Ours (wide)		Ours (deep)	
RMSNorm	LayerNorm	Type Emb.	APE $\downarrow$	JPE $\downarrow$	APE $\downarrow$	JPE $\downarrow$
	✓	✓	126.24	213.85	132.09	212.13
	✓		127.37	213.06	132.80	214.85
	✓	✓	<b>125.70</b>	<b>212.72</b>	<b>131.41</b>	<b>211.76</b>



1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108 Figure 6: Visualization of trajectory predictions ( $K = 20$ ) of our (wide) model on ETH-UCY data.  
1109 X-Y coordinate values in the plots are in meters. All model predictions are in color, ground truth  
1110 future trajectories are black dashes, and input trajectories are in gray.

1111  
1112 UMPM dataset. The results confirm their importance: removing the type embeddings degrades  
1113 performance (e.g., APE increases from 125.70 to 126.24 for the wide model), confirming they are  
1114 valuable for helping the model distinguish between pose and trajectory streams. RMSNorm also  
1115 provides a consistent, albeit minor, performance benefit over LayerNorm while being more com-  
1116 putationally efficient.

## H ADDITIONAL QUALITATIVE RESULTS

### H.1 TRAJECTORY-ONLY VISUALIZATION

1123 To provide qualitative results for the trajectory-only task, Figure 6 visualizes our model’s  $K = 20$   
1124 proposals on challenging, crowded scenes from the ETH-UCY dataset. The visualizations show our  
1125 model’s ability to capture a wide, multi-modal distribution of plausible future paths, correctly iden-  
1126 tifying diverse outcomes (e.g., turning left, turning right, or stopping) in high-uncertainty scenarios.

### H.2 FAILURE CASE ANALYSIS

1129 While our model is robust, it is not without limitations, particularly in complex multi-person scenes.  
1130 Our model treats all individuals independently, which can lead to unrealistic predictions when  
1131 agents’ motions are strongly coupled or highly unusual. Figure 7 presents qualitative failure cases  
1132 from MOCAP-UMPM. The first example showcases an intricate interaction where the blue and  
1133 green agents are turning in a circle while holding hands. Our model, processing them independently,  
fails to capture this complex, coupled motion. This also highlights that integrating explicit multi-

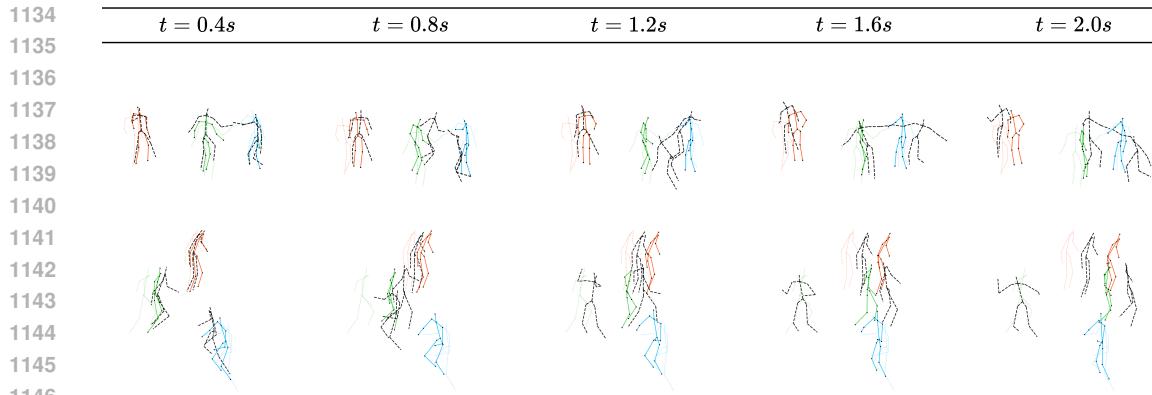


Figure 7: Visualization of predictions of our (wide) model on MOCAP-UMPM data. All model predictions are in color. Ground truth future poses are black dashes, and the last-known input positions are colored dashes.

agent interaction modules is a critical and promising direction for future work. The second example shows the blue and green agents undergoing an unexpected, rapid acceleration. The model’s predictions struggle to keep pace with this abrupt change in dynamics, likely defaulting to a smoother or more mean-reverting trajectory, and thus accumulating significant error.

## I PERFORMANCE ON HIGHLY INTERACTIVE SCENARIOS

To specifically address the model’s generalization capability on complex multi-agent interactions, we conducted an evaluation on the challenging WorldPose dataset (Jiang et al., 2024). Unlike the standard pedestrian dynamics in ETH-UCY or social mingling in UMPM, WorldPose features high-intensity sports scenarios (soccer) characterized by rapid changes in velocity, complex contact interactions, and adversarial intent. This serves as a rigorous stress test for our architecture in scenarios where prior work typically relies on dedicated interaction modules.

### I.1 EXPERIMENTAL SETUP

The WorldPose data was preprocessed to use an 80% train and 20% test split. The dataset contains player poses ( $M = 24$ ) recorded at 25 fps at soccer games. All models observed 1.0 seconds of past motion and predicted the best of  $K = 10$  proposals for 1.0 seconds into the future.

### I.2 RESULTS AND ANALYSIS

The results in Table 13 demonstrate that our simple, unified framework substantially outperforms prior complex, interaction-aware methods. Specifically, SimpliHuMoN reduces APE by 56.7% compared to T2P and 64.6% compared to EMPMP, likely due to the difficulty of modeling rapid, non-cyclic sports motions using architectures optimized for smoother walking gaits. Our method’s ability to handle this data without architectural modification highlights the universality of the proposed Transformer decoder, validating the strength of our core architecture while underscoring that integrating explicit multi-agent interaction mechanisms remains a critical and promising direction for future work.

Table 13: Comparison of APE/JPE metrics on WorldPose data. Lower values are better ( $\downarrow$ ), with the best results shown in **bold**. An asterisk (\*) denotes models we recomputed for this setup.

Model	APE $\downarrow$	JPE $\downarrow$
T2P*	362.7	913.6
EMPMP*	443.4	981.5
Ours (wide)	<b>156.8</b>	<b>746.3</b>