

IMoS: Intent-Driven Full-Body Motion Synthesis for Human-Object Interactions - Appendix

Anindita Ghosh^{1,2,3} Rishabh Dabral^{2,3} Vladislav Golyanik^{2,3} Christian Theobalt^{2,3} Philipp Slusallek^{1,3}

¹ German Research Center for Artificial Intelligence (DFKI)

²Max-Planck Institute for Informatics

³Saarland Informatics Campus

<https://vcai.mpi-inf.mpg.de/projects/IMoS>

1. Evaluation Metrics

We evaluate our method using the Mean Per Joint Positional Error (MPJPE). It measures the mean joint error over all time steps T as

$$\text{MPJPE} = \frac{1}{NT} \sum_{n \in N} \sum_{t \in T} \|J_t - \hat{J}_t\|_2, \quad (1)$$

where \hat{J} are the joint positions computed from the synthesized SMPL-X parameters and N is the size of our test set. However, our task requires that the models synthesize a diverse set of plausible motions for any type of intent. Therefore, only calculating the Euclidean error with the ground-truth motion, does not provide a complete picture of their synthesis quality.

Therefore, to understand the overall motion distribution statistics, we use the Average Variance Error [GCO*21]. The Average Variance Error (AVE) computes the L_2 error between the variance of the joint positions and that of the ground truth as

$$\text{AVE} = \frac{1}{N} \sum_{n \in N} \|\sigma - \hat{\sigma}\|_2, \text{ with } \sigma = \frac{1}{T-1} \sum_{t \in T} (J_t - \bar{J})^2, \quad (2)$$

where \bar{J} is the mean pose over T time steps, σ is the ground-truth variance, and $\hat{\sigma}$ is the variance of the synthesized sequence.

We also report four statistical metrics namely the Frechet Inception Distance (FID) [HRU*17], Recognition Accuracy, Diversity and Multimodality for a better comparison with the existing methods of Action2Motion [GZW*20] and ACTOR [PBV21].

For calculating FID, we extract features from the generated motions in our test split and the ground-truth motions in the test split, and calculate the feature distribution between the generated motions and the ground-truth motion. To extract the motion features, we train a standard RNN action recognition classifier for GRAB dataset, and use the final layer of this classifier as the motion features. A lower FID score means better quality of generated results.

Recognition accuracy indicates the correlation of the generated motions with their action types. We use the pre-trained RNN action recognition classifier to classify the motions in our test split, and calculate recognition accuracy.

Through Diversity we measure variation in the motion features across all action categories. We sample two same-sized subsets of generated motions from various action types and extract the respective set of motion features. The Diversity between these two set of motions is calculated as

$$\text{Diversity} = \frac{1}{S_d} \sum_{i=1}^{S_d} \|F_i - \hat{F}_i\|_2, \quad (3)$$

where F_1, F_2, \dots, F_{S_d} and $\hat{F}_1, \hat{F}_2, \dots, \hat{F}_{S_d}$ are the motion feature vectors of the two subsets and S_d is the sample size.

Multimodality measures how generated motion's features diversify within each action type. Given motion sequences from C different action types, for any c^{th} action, we randomly sample two subsets of same size and extract their respective motion feature vectors. Multimodality is then calculated as

$$\text{Multimodality} = \frac{1}{CS_l} \sum_{c=1}^C \sum_{i=1}^{S_l} \|F_{c,i} - \hat{F}_{c,i}\|_2, \quad (4)$$

where $F_{c,1}, F_{c,2}, \dots, F_{c,S_l}$ and $\hat{F}_{c,1}, \hat{F}_{c,2}, \dots, \hat{F}_{c,S_l}$ are the motion feature vectors of the two subsets and S_l is the sample size.

2. CLIP based embedding vs. randomly initialized vector embeddings for the intent labels

We visualize the cosine similarities of the intent vectors embedded using CLIP [RKH*] in Fig. 1. We see embeddings of intents with similar meanings such as “drink” and “pour”, “turn on” and “switch on”, “eat” and “consume”, “pass” and “transfer” have a higher cosine similarity whereas embeddings of intents with different meanings such as “offhand” and “inspect”, or “switch on” and “inspect” have low similarity values between them.

In Fig. 2, we visualize the cosine similarities between intent embeddings where the intent labels are embedded using randomly initialized vectors of 512 dimension (as done in Sec. 5.4 Ablation 1). We see that for the randomly initialized vector embeddings, there is no semantic understanding between similar or dissimilar intents.

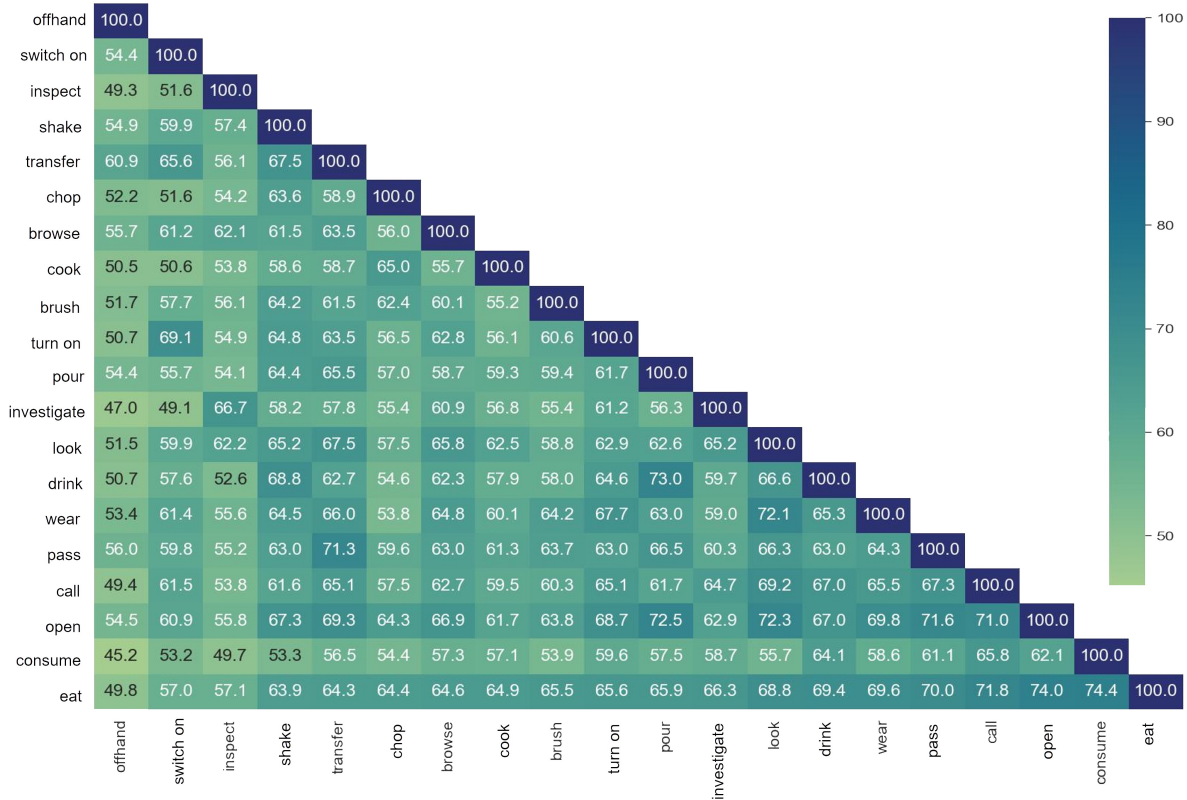


Figure 1: Confusion matrix showing the cosine similarity percentage where intent labels are encoded using CLIP [RKH*]. We see embeddings of intents with similar meanings have a higher cosine similarity whereas embeddings of intents with different meanings have low similarity values.

References

[GCO*21] GHOSH, ANINDITA, CHEEMA, NOSHABA, OGUZ, CENNET, et al. “Synthesis of Compositional Animations From Textual Descriptions”. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021 1.

[GZW*20] GUO, CHUAN, ZUO, XINXIN, WANG, SEN, et al. “Action2motion: Conditioned generation of 3d human motions”. *Proceedings of the 28th ACM International Conference on Multimedia*. 2020 1.

[HRU*17] HEUSEL, MARTIN, RAMSAUER, HUBERT, UNTERTHINER, THOMAS, et al. “Gans trained by a two time-scale update rule converge to a local nash equilibrium”. *Advances in neural information processing systems* (2017) 1.

[PBV21] PETROVICH, MATHIS, BLACK, MICHAEL J., and VAROL, GÜL. “Action-Conditioned 3D Human Motion Synthesis with Transformer VAE”. *International Conference on Computer Vision (ICCV)*. 2021 1.

[RKH*] RADFORD, ALEC, KIM, JONG WOOK, HALLACY, CHRIS, et al. “Learning Transferable Visual Models From Natural Language Supervision”. *Proceedings of the 38th International Conference on Machine Learning* 1, 2.

