

WANDR: Wrist-driven Autonomous Navigation for Data-based Goal Reaching Supplementary Material

Markos Diomataris^{1,2} Nikos Athanasiou¹ Omid Taheri¹ Xi Wang²
Otmar Hilliges² Michael J. Black¹

¹Max Planck Institute for Intelligent Systems, Tübingen, Germany ²ETH Zürich, Switzerland

1. Introduction

This supplemental material offers more details regarding the use of our method in an optimization framework, the effect of motion duration on the generated motions, different applications of our method, and more qualitative results. Please see the **Supplementary Video**, where we extensively demonstrate the realism and adaptability of our generated reaching motions across diverse scenarios.

The video effectively contains: (1) the problem and our motivation, (2) our method and key ideas, (3) multiple example motions, and (4) different applications of our method such as extension to reaching dynamic goals. The video serves as a dynamic and illustrative supplement, showcasing our contributions in a manner that is hard to show in a paper format.

2. Technical Implementation

2.1. Model Architecture

WANDR c-VAE architecture [2] employs an Encoder and a Decoder, each composed of fifteen layers in a Multi-Layer Perceptron (MLP) configuration. We integrate relu activation functions, dropout and layer normalization at each stage for enhanced performance. The latent space is represented as a 64-dimensional vector. In our design, the condition signal of the c-VAE is concatenated with the input delta (in the case of the Encoder) and the latent vector (in the case of the Decoder).

2.2. Training Details

We developed and trained our method using the PyTorch framework [3]. We train it for 900 epochs on 4 Tesla V100 GPUs. We use a batch size of 512, resulting in approximately 20 hours of training duration. For optimization, we use Adam [1] with a starting learning rate of $1e - 4$ that linearly decreases to $1e - 5$ during training.

A crucial aspect of our training regimen includes performing a teacher-forcing method, which involves feeding the model’s own predictions back into the input. This pro-

cess facilitates the Decoder network in acquiring the capability to compensate for potential errors that may arise during the prediction of deltas. During the whole process of the training, the c-VAE is being trained on the task of auto-encoding the motion deltas. As the training progresses, we additionally perform motion generation for a few steps. Specifically, we reconstruct the deltas, integrate them to obtain the subsequent pose, and then sample from the latent space while conditioning on the generated pose. We repeat this process for up to s steps increasing the s linearly from 0 up to 10 along the span of 50 epochs and then keeping it fixed.

2.3. Optimization Details

The formulation of our method as a c-VAE provides us with a smooth latent space that allows us to search this manifold in an optimization process to reach various target goals. For this, we apply specific constraints to the decoder’s output, i.e. the body poses, and optimize the latent space representation of the poses to achieve the desired motions.

In Tab. 2 of the main manuscript, we explore how optimizing a trained motion prior performs compared to our method without any optimization, in achieving various goals. The results show that WANDR without optimization performs better than the other methods, even with optimization.

To do the optimization, we first generate a motion to reach a goal using WANDR, and then refine the motion through optimization aiming to align the wrist’s position in the final frame more closely with the target goal.

In order to achieve this, we employ a dual-component loss function:

$$\mathcal{L}_{opt} = \mathcal{L}_{norm} + \mathcal{L}_{goal}$$

Here, \mathcal{L}_{norm} represents the log-likelihood of the motion’s latent vectors under a normal distribution. \mathcal{L}_{goal} calculates the mean square error between the wrist’s final frame location and the goal. \mathcal{L}_{norm} seeks to maintain the generated motion within plausible human movements, while \mathcal{L}_{goal}

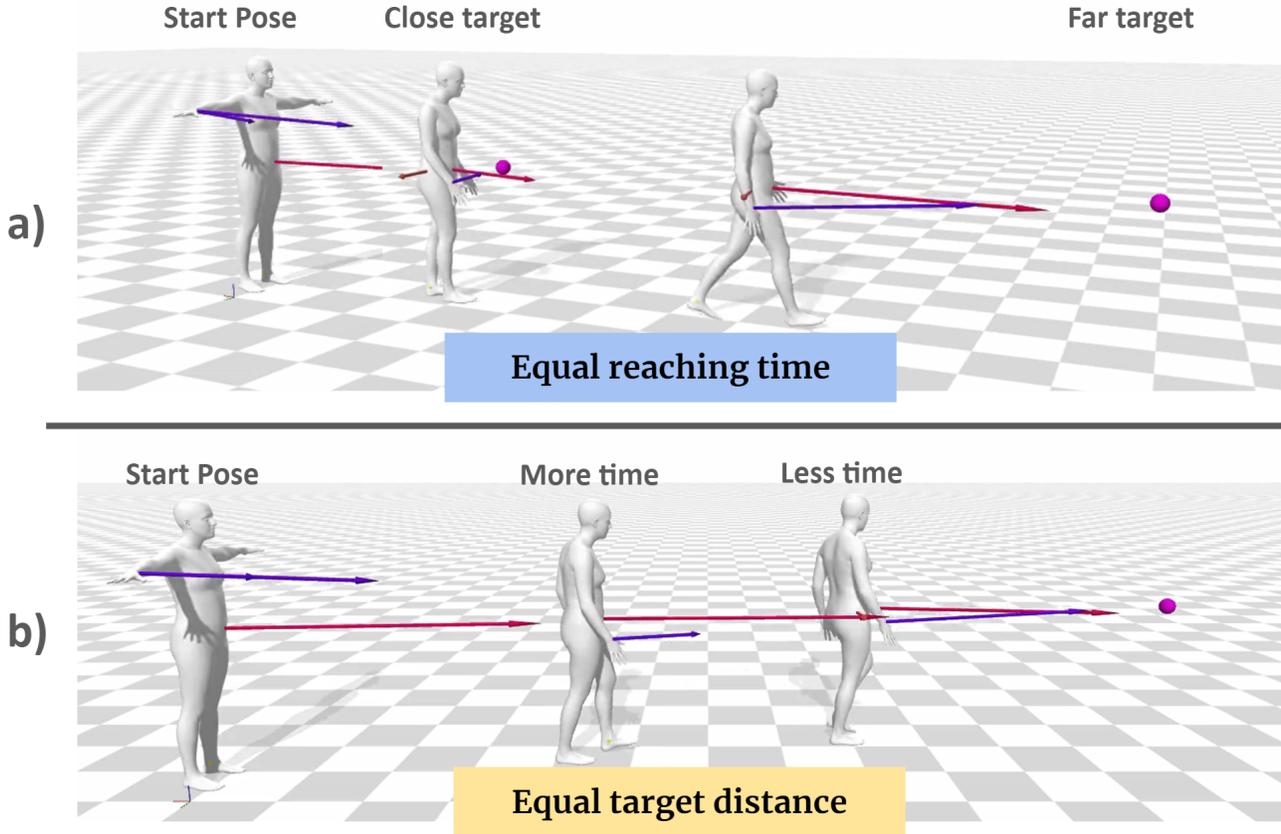


Figure S.1. Generated motions for different goals with the same time constraint (a) and for a similar goal with varying time constraints (b). The results show that our dynamic intent features enable the adaptability of WANDR to generate time-controlled motions.

specifically tunes the motion to bring the wrist in proximity to the goal in the final frame. Notably, even though \mathcal{L}_{goal} is applied only on the final frame, its gradient flows across the whole sequence since the autoregressive generation process is fully differentiable.

Sec. 3 shows our method’s ability to produce motions with different time durations, as well as its integration with the optimization framework presented in the main paper.

3. WANDR Applications

In this section we show that our method can be used in various scenarios and for different applications.

3.1. Time-controlled Motions

One key aspect of our intention features is the dependency of the wrist-intention vector on the goal-reaching time, enabling the generation of time-controlled motions. As mentioned in the main paper, this vector is computed by dividing the distance from the wrist to the goal by the time remaining to reach it. During inference, by changing the reaching time or distance, the generated motions adapt and become

rapid or slow. Figure S.1 presents two scenarios: (a) reaching different goals within the same time duration, and (b) reaching a goal within different time durations. These cases illustrate how the motion varies in response to the time and distance parameters. For qualitative examples, please see the **Supplementary Video**.

3.2. Optimization-enabled extensions

Multi-goal Reaching: We show that our unique intention features enable the generation of motions to achieve multiple goals sequentially. Although trained for single-goal achievement, the dynamic nature of our intention features allows for multiple goal definitions during inference. These features are recalculated and updated at each iteration of WANDR’s autoregressive process, adapting to changes in goal locations. Figure S.2 illustrates this capability, where a motion sequence is generated to achieve several goals. This adaptability also extends to tracking and following moving targets, as shown in our **Supplementary Video**.

Waypoint Following: Our method extends to the application of reaching goals while at the same time having the virtual human passing through arbitrary waypoints. Specif-

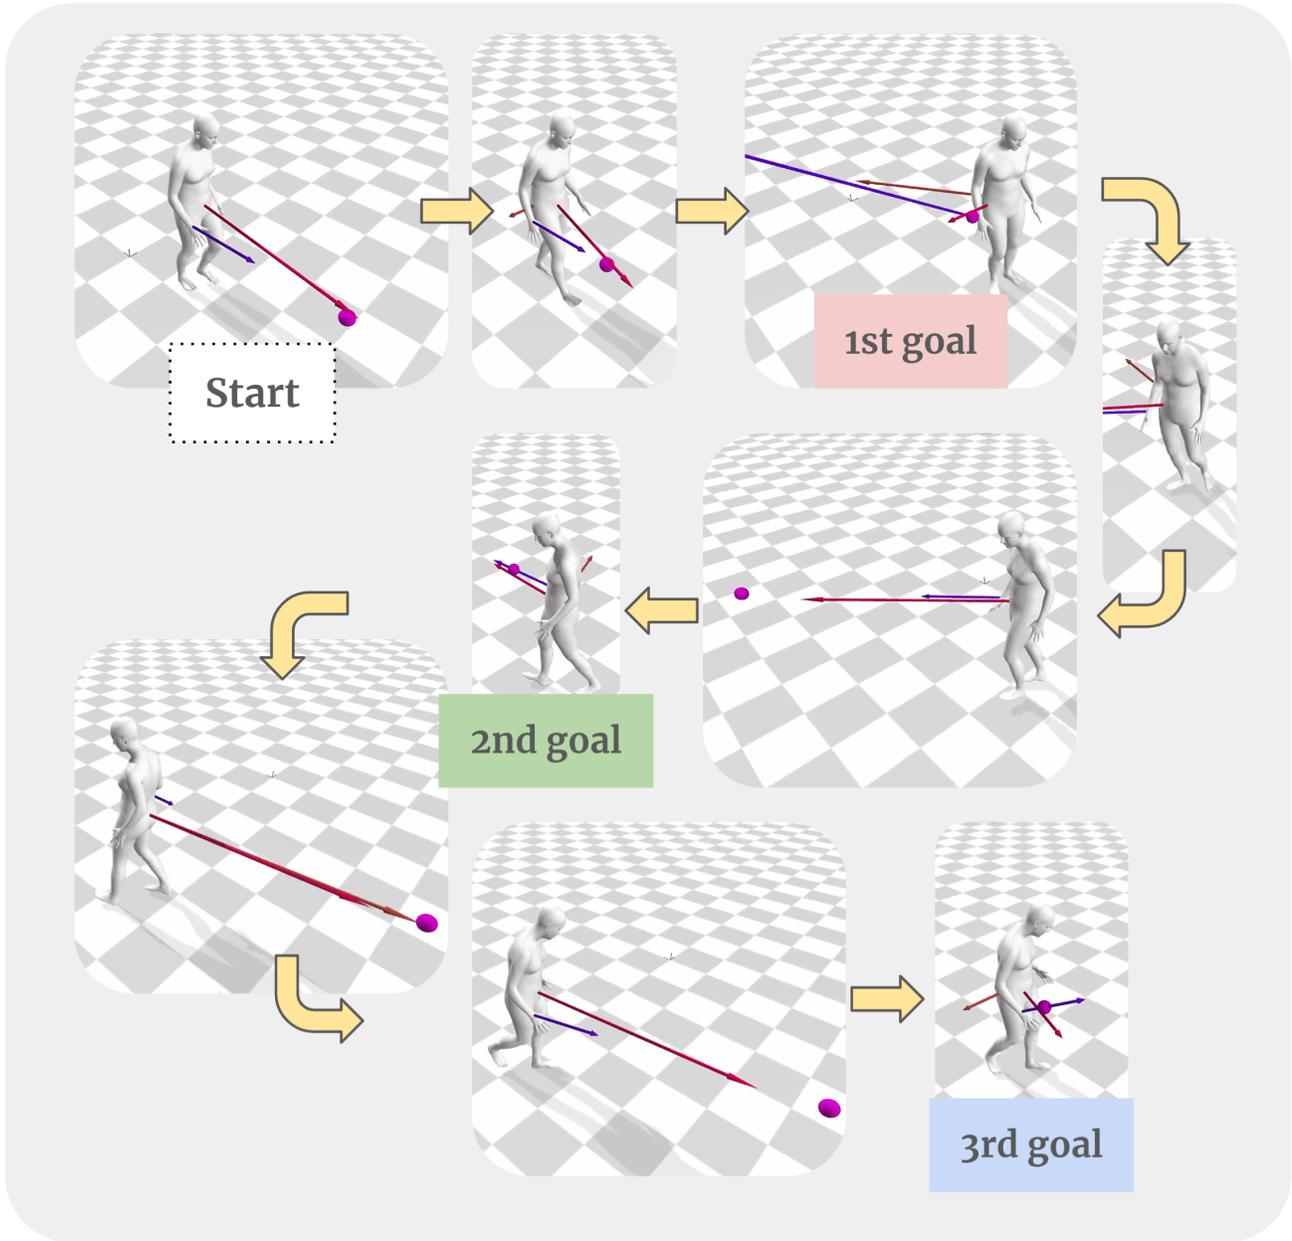


Figure S.2. A demonstration of WANDR generating a motion sequence to achieve multiple goals. The intention features are recalculated and updated after each iteration of the autoregressive process, enabling dynamic goal adjustments during the motion generation.

ically, we are able to choose a waypoint and have the human pass from it at a chosen frame, while still reaching for the goal at the end of the motion. To achieve this, we extend the optimization approach described in section Sec. 2.3 with the addition of an extra mean square error loss between the ground projection of the pelvis and the waypoint location for a particular frame. Fig. S.3 showcases an example of this application, underlining the adaptability of WANDR in navigating through waypoints while simultaneously reach-

ing a goal (please see video for an example motion). It is worth noting that passing through waypoints can be trivially extended to following trajectories, since a trajectory can be approximated by sampled waypoints on a curve.

3.3. Extending WANDR to other joints

WANDR can be, in a straight forward way, extended to other joints just by replacing the wrist position with the joint of interest in the definition of I^w . For example, by re-

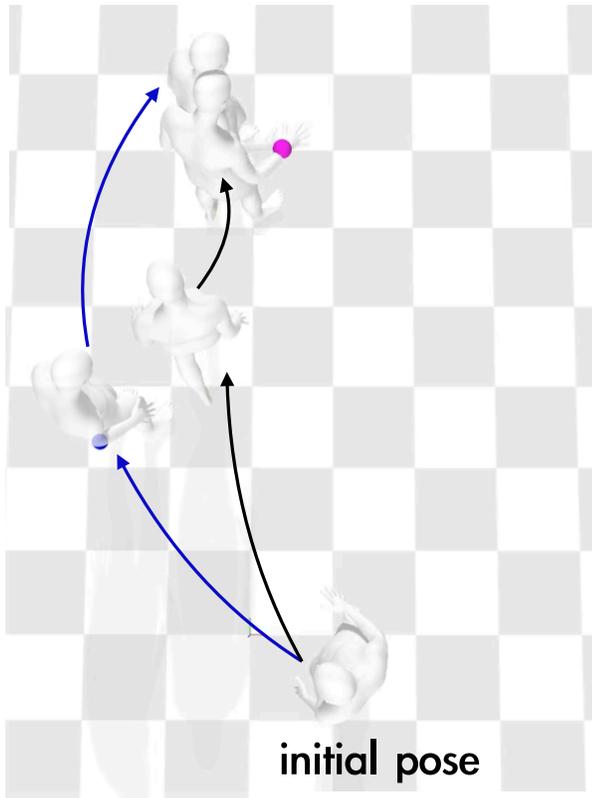


Figure S.3. A generated motion from WANDR for waypoint (blue sphere) following while reaching for a goal (pink sphere). The initial generation of WANDR follows the motion marked with the black arrows. After optimization, the motion manages to pass through the waypoint, while still reaching for the goal at the end of the motion. This illustrates that WANDR provides a smooth latent space for the motions that can be aligned with the goal-reaching motion with predefined waypoints using an optimization process.

placing the wrist with the pelvis joint, we can get a motion generator that can produce motions that follow waypoints. However, we note that controlling multiple body joints simultaneously is not trivial with the current design. It would require redesigning the intention features to enable learning which joint corresponds to which intention feature.

4. Perceptual Study

Throughout our experiments, we empirically found that the foot skating metric has a high correlation with the quality of the motion. Nevertheless, in order to properly evaluate the perceptual quality of WANDR’s generated motions we conduct two perceptual studies through amazon mechanical turk. The studies aim at quantifying how close WANDR’s motions are perceptually compared to real human motions taken from AMASS. In the first study, users rate the realism of the motions using with a 5-level Likert scale (1 → non-realistic & 5 → realistic). Only one motion is shown

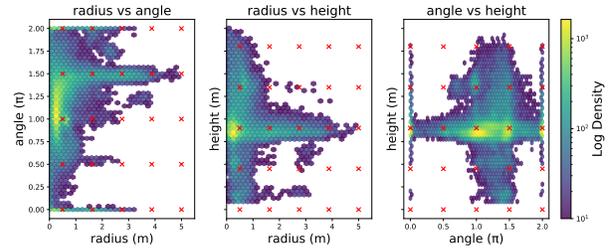


Figure S.4. Overlay of the distribution of training pseudo-goals with the evaluation goals (×) of WANDR in all pairwise combinations of their cylindrical coordinates. Our evaluation goals uniformly cover a range of goals both outside and inside the training distribution.

at a time. In the second study, users are asked to choose the most realistic motion between two, one coming from WANDR and one coming from an AMASS sequence. We clip motions to a 2 second duration and only show motions from WANDR that succeeded in reaching their goal.

In the first study, AMASS ground-truth motions score $3.8 \pm 1 / 5$ vs $3.4 \pm 1 / 5$ for WANDR. The comparative study finds that 30.2% of the users preferred WANDR motions over AMASS. These findings indicate that the WANDR motions are perceptually close to real motions.

5. Evaluation Distribution

To better demonstrate that WANDR has been evaluated on out-of-distribution data, in Fig. S.4 we visualize the density of the pseudo-goal training locations (in cylindrical coordinates) and overlay the goal locations (marked as ×) used to evaluate WANDR. We clearly observe that most evaluation goals lie on either low probability or unseen locations.

References

- [1] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 1
- [2] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 1
- [3] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 1