

# Controlling Intent Expressiveness in Robot Motion with Diffusion Models

Wenli Shi, Clemence Grislain, Olivier Sigaud, Mohamed Chetouani

Sorbonne Université, CNRS, Institut des Systèmes Intelligents et de Robotique (ISIR)  
Paris, France

{wenli, grislain, sigaud, chetouani}@isir.upmc.fr

**Abstract**—Legibility of robot motion is critical in human–robot interaction, as it allows humans to quickly infer a robot’s intended goal. Although traditional trajectory generation methods typically prioritize efficiency, they often fail to make the robot’s intentions clear to humans. Meanwhile, existing approaches to legible motion usually produce only a single “most legible” trajectory, overlooking the need to modulate intent expressiveness in different contexts. In this work, we propose a novel motion generation framework that enables controllable legibility across the full spectrum—from highly legible to highly ambiguous motions. We introduce a modeling approach based on an Information Potential Field to assign continuous legibility scores to trajectories, and build upon it with a two-stage diffusion framework that first generates paths at specified legibility levels and then translates them into executable robot actions. Experiments in both 2D and 3D reaching tasks demonstrate that our approach produces diverse and controllable motions with varying degrees of legibility, while achieving performance comparable to SOTA. Code and project page: <https://legibility-modulator.github.io/>.

## I. INTRODUCTION

When collaborating with humans, robots must not only reach their goals efficiently but also move in ways that make their intentions clear. A trajectory that minimizes path length may be efficient, but when multiple possible goals exist, it can appear ambiguous to an observer. In contrast, a legible trajectory enables human observers to quickly and confidently infer the true goal of the robot [1]. Thus, legibility is critical for fluent and safe interaction.

Various methods have been proposed to generate legible robot motion, ranging from optimization-based planners [2], [3] to learning-driven generative approaches [4]. These works show that trajectories can be designed to make a robot’s goal more transparent to humans. However, most existing approaches focus on producing a single maximally legible trajectory. This overlooks the fact that legibility is not binary, but a spectrum: in cooperative tasks, robots may need to be highly transparent; in time-critical tasks, they may prioritize efficiency; and in adversarial contexts, they may even benefit from intentional ambiguity. Current methods do not provide a principled mechanism for modulating this degree of intent expressiveness.

To address this need, we propose a framework for generating trajectories with controllable levels of legibility. Our key idea is to introduce an Information Potential Field (IPF) that models the probabilistic structure of goal inference and yields a potential-aware legibility score. This score enables smooth modulation across the legibility spectrum.

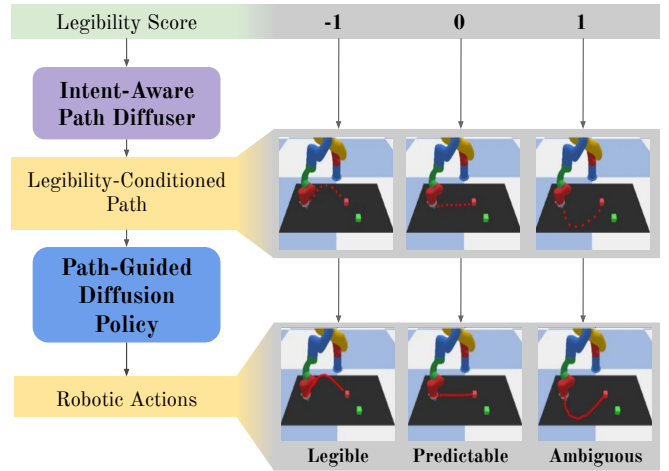


Fig. 1. **Overview of Legibility Modulator.** The *Intent-Aware Path Diffuser* generates an intermediate path conditioned on the desired intent expressiveness (legibility score). The *Path-Guided Diffusion Policy* then translates this path into executable robot actions, producing motion that aligns with the specified level of legibility.

To support this approach, we first construct a diverse trajectory dataset inspired by the Quality-Diversity paradigm [5]. This dataset covers a wide range of deviation patterns that are retrospectively labeled with potential-based legibility scores, ensuring broad coverage of the legibility space. Building on this foundation, we propose **Legibility Modulator**, a two-stage diffusion framework illustrated in Figure 1.

We validate our method in 2D and 3D reaching tasks with various goal configurations. Results show that Legibility Modulator produces diverse and controllable trajectories across the entire legibility spectrum, while maintaining task success and matching the performance of state-of-the-art baselines in maximally legible settings.

Our contributions are threefold:

- 1) We introduce a new legibility-score based on an information potential field, enabling continuous scoring of trajectories in terms of intent expressiveness.
- 2) We construct a trajectory dataset using a quality-diversity paradigm to generate diverse trajectories across the legibility spectrum.
- 3) We propose a two-stage diffusion approach trained on this dataset, achieving performance comparable to SOTA methods while additionally enabling fine-grained control of intent expressiveness from ambiguous to highly legible behaviors.

## II. RELATED WORK

**Legible Robot Motion:** Legible motions enable human observers to quickly and confidently infer the intended goal of a robot, which in most studies corresponds to a spatial target in the workspace. Dragan et al. [1] provided a probabilistic formulation distinguishing legibility from predictability, highlighting its importance for safe and fluent human-robot interaction [6]. Early works employed optimization-based planners that shaped cost functions on trajectory characteristics such as distance to alternative goals [7]. Although effective, these approaches required hand-crafted heuristics and lacked scalability. Learning-based methods have since emerged [8], [9], including reinforcement learning to optimize legibility metrics [10] and supervised observer models to predict human goal inference [11].

**Robot Motion from Conditional Generative Models:** Conditional generative models have recently shown strong capabilities in robot motion generation [12]. They naturally capture multi-modality and provide stable training. In particular, diffusion models [13] have been applied to both action- and trajectory-level generation [14], [4], offering temporal consistency and robustness to multi-modality. Methods such as classifier-free guidance [15] further allow the sampling process to be biased toward conditional distributions in a controllable manner. Building on these, [4] proposed Legibility Diffuser, a diffusion-based policy that leverages guided sampling to imitate the most legible modes from demonstrations. However, such approaches focus on producing a single maximally legible trajectory. Since legibility is inherently a spectrum that must be modulated depending on the context, this limitation motivates our work.

**Potential Fields:** Artificial Potential Fields (APF), introduced by Khatib [16], have been widely used for navigation and obstacle avoidance. Beyond basic planning, they have also been combined with generative models. Luo et al. [3] proposed potential-based diffusion planning to avoid local minima, while [17] applied entropy-scaled fields to improve legibility in cluttered environments. However, prior work treats potential fields as geometric constructs that ensure feasibility, without accounting for the probabilistic nature of human goal inference. In contrast, our work extends this paradigm by introducing a probabilistic potential formulation tailored to legibility-aware motion generation.

## III. PROBLEM STATEMENT

We formalize the problem of legible motion generation as a goal-conditioned Markov Decision Process (MDP), defined as:  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \rho_0, \mathcal{G})$ , where  $\mathcal{S}$  is the state space, typically the robot's configuration or the pose of the end effector,  $\mathcal{A}$  is the action space,  $\mathcal{T}(s'|s, a)$  is the transition function,  $\rho_0$  is the initial state distribution and  $\mathcal{G}$  is the finite set of target positions. In this framework, we define a *trajectory* as a sequence of state-action pairs  $\xi = (s_0, a_0, s_1, a_1, \dots, s_T)$ , and a *path*, a subsampled sequence of states  $\tau = (s_n, s_{2n}, \dots, s_{kn})$  where  $n$  is the step size controlling the temporal resolution. Given a start state  $s_0$  and an intended goal  $g^* \in \mathcal{G}$ , the robot must generate a

trajectory  $\xi$  that reaches  $g^*$  while maximizing the human observer's ability to infer this goal early and unambiguously, distinguishing it from alternative non-goals  $g^- \in \mathcal{G} \setminus \{g^*\}$ .

## IV. PREDICTABILITY AND LEGIBILITY

In multi-goal settings, robot trajectories must balance efficiency with intent expressiveness. Two key but distinct notions are *predictability* and *legibility* [1], [6].

**Predictability.** Given a known goal  $g$ , predictability measures how well a trajectory aligns with an observer's expectation, typically some minimal cost  $C(\xi)$  (e.g., path length). In other words, the most predictable path is simply the lowest-cost trajectory from the start to  $g$ .

**Legibility.** Legibility reverses this reasoning: from a partial trajectory  $\xi_{s \rightarrow q}$ , the observer infers the goal that maximizes the posterior probability  $P(g|\xi_{s \rightarrow q})$ . A trajectory is more legible if this probability mass quickly concentrates on the true goal  $g^*$ . More generally, the legibility of  $\xi$  can be quantified as [1]:

$$\text{Legibility}(\xi) = \frac{\int_0^T P(g^*|\xi_{[0 \rightarrow t]})f(t) dt}{\int_0^T f(t) dt}, \quad (1)$$

where  $f(t)$  emphasizes early cues and  $T$  is the time horizon.

In a reaching task with two adjacent targets, a straight line is efficient and thus predictable but ambiguous, whereas an early deviating curve is less efficient yet more legible.

## V. METHODS

Our goal is to generate robot motions that are both goal-directed and legible to human observers, with controllable levels of legibility. To this end, we propose **Legibility Modulator**, illustrated in Figure 2.

### A. Information Potential Field for Legibility Modeling

In the formal definition of legibility (Eq. (1)), the posterior  $P(g^*|\xi_{[0 \rightarrow t]})$  reflects how observers infer goals from motion prefixes, but is impractical to obtain directly. We therefore approximate it with an *Information Potential Field* (IPF), which considers spatial configurations but omits dynamic aspects like velocity.

The likelihood of observing configuration  $x$  when aiming for  $g$  is modeled as  $P(x|g) \sim \mathcal{N}(x_g, \sigma^2 I)$ .

The posterior over the intended goal  $g^*$  is then

$$P(g^*|x) = \frac{P(x|g^*)}{P(x|g^*) + \sum_{g^- \in \mathcal{G} \setminus \{g^*\}} P(x|g^-)}. \quad (2)$$

This posterior naturally encodes the competition between the intended goal and opposing ones: evidence for  $g^*$  increases  $P(g^*|x)$ , while support for distractors reduces it.

We define the *information potential* as the negative log-posterior probability of goal  $g^*$  given observation  $x$ :

$$\phi(x|g^*) = -\log P(g^*|x), \quad (3)$$

which is small when  $x$  supports  $g^*$  and large when ambiguous. This potential provides a continuous, tractable surrogate for goal disambiguation along a trajectory. For practical representation, the IPF can be stored as a spatial grid: in

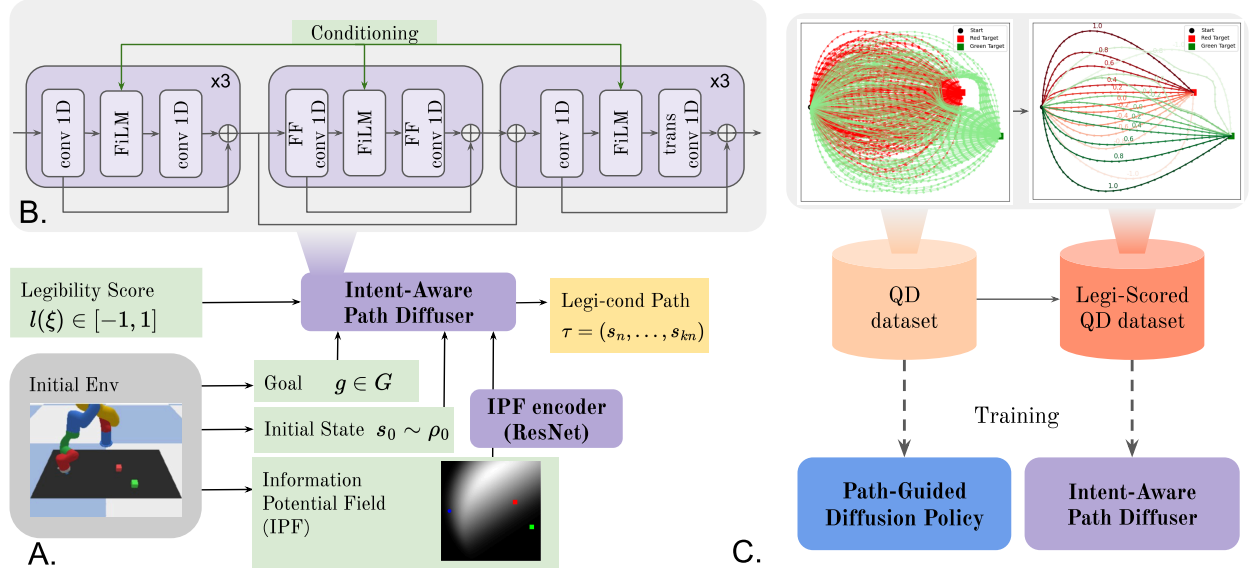


Fig. 2. **Detailed architecture of our two-stage conditional diffusion framework.** (A) The framework takes as input the initial environment state  $s_0$ , a candidate goal  $g$ , the IPF that encodes the probabilistic structure of goal inference, and a continuous *legibility score*  $\ell(\xi)$  specifying the desired level of expressiveness. (B) The Intent-Aware Path Diffuser is a U-Net conditioned via FiLM layers on the inputs, which generates a path  $\tau$  using diffusion process. (C) The Path-Guided Diffusion Policy is trained on trajectories generated by the QD algorithm, using the corresponding paths as conditioning inputs. The Intent-Aware Path Diffuser is trained on the same QD trajectories, labeled with potential-based legibility scores.

2D tasks as a single-channel image  $1 \times H \times W$ , and in 3D tasks as a volumetric grid  $D \times H \times W$ .

### B. Trajectory Dataset Generation

To enable trajectory generation across the full spectrum of legibility, we construct a dataset following the quality-diversity (QD) paradigm [5], [18]. The goal is to obtain trajectories that are both feasible and well spread in terms of their intent expressiveness, covering legible, predictable, and ambiguous behaviors.

Each trajectory is parameterized as a cubic Bézier curve [19] between the start  $s$  and goal  $g^*$ , with control points  $c_1, c_2$  sampled within bounds. Curves are discretized by arc-length to ensure consistent resolution. To characterize diversity, we extract the *most deviating point* from the straight line  $sg^*$  and use its spatial position as a descriptor. By uniformly partitioning the workspace into grid cells based on this deviation point, we obtain a set of trajectories that cover different geometric styles and legibility levels.

Each discretized trajectory  $\xi = \{x_t\}_{t=1}^T$  is further assigned a potential-based legibility score:

$$L_p(\xi) = - \sum_{t=1}^T f(t) \phi(x|g^*), \quad f(t) = \exp(-\alpha t), \quad (4)$$

where earlier points are weighted more strongly. Larger  $L_p(\xi)$  corresponds to higher legibility. We normalize the scores into  $[-1, 1]$  by ranking within each target set using,  $\ell(\xi) = 2 \cdot \frac{\text{rank}(L_p(\xi))}{N} - 1$ , where  $\ell(\xi) = 1$  denotes maximally legible,  $\ell(\xi) = -1$  maximally ambiguous, and  $\ell(\xi) = 0$  maximally predictable trajectories. This procedure yields a dataset that is both geometrically diverse and uniformly

distributed along the legibility spectrum, forming a principled basis for conditional generation.

### C. Two-Stage Conditional Diffusion Framework

**Stage 1: Intent-aware Path Diffuser:** Since legibility depends on global trajectory shape rather than low-level actions, the first stage generates legibility-conditioned paths  $\tau = \{s_n, \dots, s_{kn}\}$  (we set  $k = 8$  for compactness). The diffusion model is conditioned on three signals: (i) start and goal  $(s, g)$ , (ii) Information Potential Field (IPF) features, and (iii) the legibility label  $\ell(\xi)$ . IPF features are encoded as images in both 2D and 3D environments, enabling integration with convolutional backbones. These inputs are fused into a context vector and injected into a UNet denoiser via FiLM [20], which refines Gaussian noise into a legibility-consistent trajectory plan.

**Stage 2: Path-Guided Diffusion Policy:** The second stage converts the path into executable actions. It follows the diffusion policy framework [14] and is conditioned on: (i) a short state history, (ii) the goal  $g$ , and (iii) the path  $\tau$ . The policy outputs an action sequence  $(a_1, \dots, a_H)$  that reproduces the spatial plan while satisfying dynamics.

Both models are trained independently - Stage 1 on trajectories labeled with legibility scores, and Stage 2 on state-action demonstrations - and are composed sequentially at inference ensuring control of legibility at the trajectory level and feasibility.

## VI. EXPERIMENTS

We compare our **Legibility Modulator** against three baselines: **Legibility Diffuser** [4], a diffusion-based policy

with guided sampling toward legible modes, **Diffusion Policy** [14], a standard conditional diffusion baseline without legibility guidance, and **Dataset Oracle**, the most legible trajectories from demonstrations, serving as an upper bound.

We evaluate them on the *Block Reaching* task, a standard benchmark for legible motion [1], [11], where a robot must reach one of two possible blocks. We consider two variants: (i) a 2D setting modeling end-effector motion in a plane, adapted from prior multi-modal policy tasks [21], [22], and (ii) a 3D setting in PandaGym [23] with a Franka arm, closer to real manipulation. We report two complementary legibility metrics.

- **Distance-based legibility score** used in [4] measures geometric divergence from distractor goals:

$$L_d(\xi) = \sum_{t=1}^T \frac{\|g^- - x_t\|^2}{t} \quad (5)$$

where  $g^-$  is the non-target goal. Higher values indicate clearer separation, though the metric is less sensitive to subtle legibility differences.

- **Potential-based legibility score** uses our IPF measure in Eq. (4) with  $f(t) = -1/t$ , for a fair comparison with  $L_d$ . Lower values indicates more legible motions.

## VII. RESULTS

### A. Comparison in the Maximally Legible Regime

Most prior methods on legible motion generation focus on producing a single maximally legible trajectory. To align with this setting, we first fix the conditioning label to  $\ell(\xi) = 1$  and compare against the baselines. Table I shows that our method matches the performance of the Legibility Diffuser on both  $L_d$  and  $L_p$ , while maintaining perfect task success. Furthermore, it achieves significant improvements in  $L_d$  and  $L_p$  for both 2D and 3D environments compared to the legibility-unaware Diffusion Policy. These results indicate that our framework produces highly legible motion predictions that align with SOTA legibility-aware methods when evaluated under the maximally legible regime established in the literature.

### B. Controllable Legibility Spectrum

Beyond the maximally legible setting, Legibility Modulator uniquely enables fine-grained control of legibility. By varying the conditioning label  $\ell(\xi)$  from  $-1$  (most ambiguous) through  $0$  (predictable) to  $1$  (most legible), trajectories evolve smoothly and qualitatively align with the intended expressiveness. Figure 3 illustrates this effect in the 3D reaching task. On the left, we show example trajectories generated by our method, colored from light to dark red to indicate increasing expected legibility scores. When conditioned on a low legibility score (lighter red), Legibility Modulator produces ambiguous trajectories with respect to the true goal (red). In contrast, when conditioned on a high legibility score (darker red), it produces highly legible trajectories from which a human observer can better infer

TABLE I

COMPARISON OF OUR METHOD UNDER THE MAXIMALLY LEGIBLE REGIME  $\ell(\xi) = 1$  WITH BASELINES, ON THE 2D AND 3D BLOCK REACHING ENVIRONMENTS. IN BOLD ARE THE RESULTS THAT ARE SIGNIFICANTLY DIFFERENT FROM LEGIBILITY-UNAWARE METHODS.

Method (2D Env)	SR	Target R		Target G	
		$L_d \uparrow$	$L_p \downarrow$	$L_d \uparrow$	$L_p \downarrow$
Diffusion Policy	1.	1.54	$1.5e^{-3}$	1.23	$2.52e^{-1}$
Legibility Diffuser	1.	<b>1.97</b>	<b><math>3e^{-4}</math></b>	<b>1.67</b>	<b><math>4.62e^{-2}</math></b>
Legi Modulator (Ours)	1.	<b>1.96</b>	<b><math>4e^{-4}</math></b>	<b>1.65</b>	<b><math>4.19e^{-2}</math></b>
Oracle (Max Leg Traj)	-	2.02	$3e^{-4}$	1.69	$3.24e^{-2}$
<b>Method (3D Env)</b>					
Diffusion Policy	1.	3.05	$3.01e^{-1}$	3.24	$2.80e^{-2}$
Legibility Diffuser	1.	<b>3.85</b>	<b><math>3.05e^{-2}</math></b>	<b>3.95</b>	<b><math>3.30e^{-3}</math></b>
Legi Modulator (Ours)	1.	<b>3.88</b>	<b><math>2.18e^{-2}</math></b>	<b>3.95</b>	<b><math>1.50e^{-3}</math></b>
Oracle (Max Leg Traj)	-	3.94	$1.24e^{-2}$	3.74	$3.10e^{-3}$

the intended goal. On the right, we show that the potential-based score  $L_p$  decreases smoothly and monotonically as  $\ell(\xi)$  increases, demonstrating that our method achieves fine-grained modulation across the legibility spectrum.




Fig. 3. Results in the 3D Block Reaching (Target Red), showing smooth modulation from ambiguous to legible motions. Similar behaviors are also observed for the alternative target and in the 2D environment.

In summary, Legibility Modulator not only matches existing methods under maximally legible settings but also provides a principled mechanism for modulating intent expressiveness across the full spectrum, while preserving task success.

## VIII. CONCLUSION

We presented Legibility Modulator, a two-stage diffusion framework for generating robot motion with controllable levels of legibility, combining an Information Potential Field for trajectory scoring with an intent-aware path diffuser and a path-guided policy. Experiments in 2D and 3D reaching tasks showed that our approach not only matches SOTA performance in maximally legible settings but also enables smooth modulation across the entire legibility spectrum. A limitation of this study is that experiments were restricted to simulated reaching tasks, which do not fully capture the complexity of real human–robot interaction. In future work, we plan to extend our framework to more naturalistic settings and evaluate its effectiveness with human participants.

## ACKNOWLEDGMENTS

This work was performed using HPC resources from GENCI–IDRIS (Grant 2025-[AD011016424]). This research was partially funded by the French National Research Agency (ANR) under the OSTENSIVE project “ANR-24-CE33-6907-01.”

## REFERENCES

- [1] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, “Legibility and predictability of robot motion,” in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 301–308.
- [2] J. Amirian, M. Abrini, and M. Chetouani, “Legibot: Generating legible motions for service robots using cost-based local planners,” in *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*. IEEE, 2024, pp. 461–468.
- [3] Y. Luo, C. Sun, J. B. Tenenbaum, and Y. Du, “Potential based diffusion motion planning,” *arXiv preprint arXiv:2407.06169*, 2024.
- [4] M. Bronars, S. Cheng, and D. Xu, “Legibility diffuser: Offline imitation for intent expressive motion,” *IEEE Robotics and Automation Letters*, 2024.
- [5] J.-B. Mouret and J. Clune, “Illuminating search spaces by mapping elites,” 2015. [Online]. Available: <https://arxiv.org/abs/1504.04909>
- [6] T. Chakraborti, A. Kulkarni, S. Sreedharan, D. E. Smith, and S. Kambhampati, “Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior,” in *Proceedings of the international conference on automated planning and scheduling*, vol. 29, 2019, pp. 86–96.
- [7] C. Bodden, D. Rakita, B. Mutlu, and M. Gleicher, “A flexible optimization-based method for synthesizing intent-expressive robot arm motion,” *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1376–1394, 2018.
- [8] X. Zhao, T. Fan, D. Wang, Z. Hu, T. Han, and J. Pan, “An actor-critic approach for legible robot motion planner,” in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 5949–5955.
- [9] B. Busch, J. Grizou, M. Lopes, and F. Stulp, “Learning legible motion from human–robot interactions,” *International Journal of Social Robotics*, vol. 9, no. 5, pp. 765–779, 2017.
- [10] M. Bied and M. Chetouani, “Integrating an observer in interactive reinforcement learning to learn legible trajectories,” in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2020, pp. 760–767.
- [11] S. Wallkötter, M. Chetouani, and G. Castellano, “Slot-v: supervised learning of observer models for legible robot motion planning in manipulation,” in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2022, pp. 1421–1428.
- [12] R. Wolf, Y. Shi, S. Liu, and R. Rayyes, “Diffusion models for robotic manipulation: A survey,” *arXiv preprint arXiv:2504.08438*, 2025.
- [13] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [14] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, 2024.
- [15] J. Ho and T. Salimans, “Classifier-free diffusion guidance,” *arXiv preprint arXiv:2207.12598*, 2022.
- [16] O. Khatib, “Real-time obstacle avoidance for manipulators and mobile robots,” *The international journal of robotics research*, vol. 5, no. 1, pp. 90–98, 1986.
- [17] M. Schmidt-Wolf, T. Becker, D. Oliva, M. Nicolescu, and D. Feil-Seifer, “Through the clutter: Exploring the impact of complex environments on the legibility of robot motion,” *arXiv preprint arXiv:2406.00119*, 2024.
- [18] J. Huber, F. Hélienon, H. Watrelot, F. B. Amar, and S. Doncieux, “Domain randomization for sim2real transfer of automatically generated grasping datasets,” in *2024 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2024, pp. 4112–4118.
- [19] J.-H. Hwang, R. C. Arkin, and D.-S. Kwon, “Mobile robots at your fingertip: Bezier curve on-line trajectory generation for supervisory control,” in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 2. IEEE, 2003, pp. 1444–1449.
- [20] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, “Film: Visual reasoning with a general conditioning layer,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [21] N. M. M. Shafiullah, Z. J. Cui, A. Altanzaya, and L. Pinto, “Behavior transformers: Cloning  $k$  modes with one stone,” in *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. [Online]. Available: <https://openreview.net/forum?id=agTr-vRQsa>
- [22] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” *Conference on Robot Learning (CoRL)*, November 2021.
- [23] Q. Gallouédec, N. Cazin, E. Dellandréa, and L. Chen, “panda-gym: Open-Source Goal-Conditioned Environments for Robotic Learning,” *4th Robot Learning Workshop: Self-Supervised and Lifelong Learning at NeurIPS*, 2021.