SPADE: Sketch-guided Path Planning Augmented with Diffusion Experts

Extended Abstract

Charbel Abi Hana
IDEALworks GmbH
Munich, Germany
charbel.abihana@idealworks.com

Tatiana Ghantous
IMT Atlantique
Brest, France
tatiana.ghantous@imt-atlantique.net

Mikael Khalil
IMT Atlantique
Brest, France
mikael.khalil@imt-atlantique.net

Anthony Rizk
IDEALworks GmbH & Saint Joseph University of Beirut
Munich, Germany
anthony.rizk@idealworks.com

ABSTRACT

Incorporating human preferences into path planning for Autonomous Mobile Robots typically requires complex reward engineering or costly teleoperation. Recent imitation-learning frameworks such as SKIPP let operators sketch desired paths, but they suffer from limited generalization to unseen environments and fragile data-collection tools. We introduce SPADE, a framework that addresses both problems through two contributions: (i) an open-source ROS 2-based annotation tool for robust demonstration collection, and (ii) a novel Conditional Diffusion-augmented Behavioral Cloning (Cond-DBC) training strategy. In Cond-DBC, an image-conditioned diffusion model, applied via FiLM layers, serves as an offline expert that guides a compact U-Net policy during training by providing a margin loss over the path channel. On a 22 000-instance dataset spanning ten maps, a medium Cond-DBC model (1.9 M parameters) achieves 39.1% lower Absolute Pose Error and 33.5% lower Fréchet Inception Distance than the large SKIPP baseline (31 M parameters), while preserving real-time, on-edge inference.

KEYWORDS

Local Path Planning; Autonomous Mobile Robots; Diffusion Models; Learning from Demonstrations; Behavioral Cloning

ACM Reference Format:
Charbel Abi Hana, Tatiana Ghantous, Mikael Khalil, and Anthony Rizk. 2026. SPADE: Sketch-guided Path Planning Augmented with Diffusion Experts: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/RIHP6974

1 INTRODUCTION

Autonomous Mobile Robots (AMRs) increasingly need to follow operator-specific navigation behaviors that are difficult to encode as cost functions for classical planners [4, 6]. Imitation learning (IL) trains policies from expert demonstrations [15], yet standard behavioral cloning (BC) [13] generalizes poorly to unseen states [11]. Prior work learns navigation behaviors from demonstrations [8, 12] or iterative parameter tuning [14], but requires complex reward inference or repeated human feedback. SKIPP [9] lets operators sketch desired paths on occupancy maps and trains a U-Net [10] via BC, yet its generalization to new maps remained limited and its annotation tool relied on a deprecated SDK.

This creates a generalization–efficiency gap: diffusion models [1, 3] achieve strong generalization but are too slow for on-edge deployment (18–45 s), while compact BC models are fast but fail on unseen maps. DBC [2] partially bridges this gap by using a diffusion expert during training, but models the unconditional joint $p(a, s)$ where a and s are the action and state pairs respectively, diluting guidance. We introduce SPADE, which contributes:

- An open-source ROS 2 annotation tool for reproducible demonstration collection, replacing SKIPP's deprecated pipeline.
- Cond-DBC: a conditional diffusion-augmented BC strategy where FiLM-conditioned [7] diffusion models learn $p(a|s)$ and provide environment-aware guidance to lightweight policies.

2 PROBLEM FORMULATION

We formulate sketch-guided path planning as an image-to-image BC problem. The state s consists of three 128×128 binary images encoding the local occupancy grid, start position, and goal position. The action a is a 128×128 binary path image. A policy $\pi_\theta(a|s)$ is trained to minimize the pixel-wise binary cross-entropy $\mathcal{L}_{\text{BC}}$ between predicted paths $\hat{a}$ and expert demonstrations a. Our method augments this objective with a diffusion-based margin loss that measures whether $\hat{a}$ is consistent with the expert distribution, conditioned on the environment state s.

3 APPROACH

Conditional Diffusion Expert. We train an image-conditioned DDPM [5] that learns $p(a|s)$. The model receives s through FiLM layers [7] and denoises only the path channel, unlike standard DBC [2] which jointly models $p(a, s)$ over all four channels. This is critical: a joint model must disentangle whether low probability stems from an unusual state or an inappropriate action, diluting

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/RIHP6974

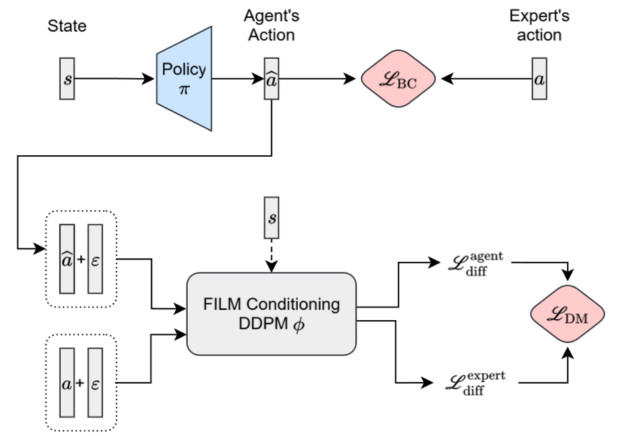

Figure 1: Cond-DBC pipeline. The BC model π predicts $\hat{a}$ from s; both $\hat{a}$ and the expert a are noised and evaluated by the FiLM-conditioned diffusion model ϕ (conditioned on s). The margin loss $\mathcal{L}_{DM}$ supplements the BC loss $\mathcal{L}_{BC}$.

guidance. By conditioning on s, Cond-DBC directly evaluates path appropriateness *for the current environment*, providing more informative gradients, particularly when the BC policy proposes paths conflicting with obstacles.

Cond-DBC Training. As illustrated in Figure 1, the frozen diffusion expert provides a margin loss $\mathcal{L}_{DM}$ during BC training. For a state s, expert action a, and predicted action $\hat{a} \sim \pi_\theta(s)$, we noise both a and $\hat{a}$ to level n and compute:

$$\mathcal{L}_{DM} = \mathbb{E}_{(s,a)\sim D}\left[\max\left(\|\hat{\epsilon}_\phi(s,\hat{a},n)-\epsilon\|^2 - \|\hat{\epsilon}_\phi(s,a,n)-\epsilon\|^2,\ 0\right)\right] \tag{1}$$

where $\hat{\epsilon}_\phi$ is the noise predicted by frozen model ϕ and ϵ is ground-truth noise. The total objective $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BC}} + \lambda\,\mathcal{L}_{DM}$ balances imitation fidelity and diffusion guidance, penalizing paths that deviate from the expert distribution with environment-aware conditioning on s.

4 EXPERIMENTAL EVALUATION

Setup. Domain experts collected 20 000 demonstrations (10 000 L-shape and 10 000 U-shape) across ten occupancy maps using our annotation tool, with 2 000 held-out instances on an unseen industrial map. These trajectory families, selected with warehouse robotics practitioners for dolly docking and last-mile delivery, represent critical operational patterns; the formulation is geometry-agnostic and extends to other shapes. We compare BC (SKIPP baseline), DBC (unconditional diffusion), and Cond-DBC at three model parameter scales: Small (117 K), Medium (1.9 M), and Large (31 M params). **Metrics**. FID (Fréchet Inception Distance, measures distributional similarity of generated paths), APE (Absolute Pose Error, measures positional accuracy in cm), and Hd_{19} (count of samples with Hausdorff distance > 19, indicates visual artifacts).

Results. Table 1 shows consistent improvements from diffusion-augmented training across all scales. The medium Cond-DBC model attains 39.1% lower APE and 33.5% lower FID than the *large* BC baseline while using 93.8% fewer parameters. Cond-DBC also avoids the artifact inflation that DBC causes in small models (DBC: 62–124 Hd_{19} vs. Cond-DBC: 43–52), confirming that modeling $p(a|s)$ yields more targeted gradients than $p(a,s)$ jointly. Ablation studies in the full version of the paper further show that when fine-tuning from L-shape to U-shape on a new industrial map, the full SPADE pipeline achieves 37.1% better FID and 62% fewer artifacts compared to baseline fine-tuning, demonstrating effective cross-behavior transfer.

Table 1: Benchmark comparison across model sizes and training approaches. Metrics: FID (lower is better), APE in cm (lower is better), Hd_{19} artifact count (lower is better).

Size	Approach	FID ↓	APE ↓	Hd_{19} ↓
		L-Shape		
Small	BC	42.27	4.43	**28**
	DBC	24.69	**2.92**	62
	Cond-DBC	**23.87**	3.14	43
Medium	BC	31.25	3.48	6
	DBC	16.51	2.21	**1**
	Cond-DBC	**13.58**	**1.46**	2
Large	BC	20.45	2.41	5
	DBC	15.72	1.89	**0**
	Cond-DBC	**12.47**	**1.47**	5
		U-Shape		
Small	BC	35.51	6.79	**52**
	DBC	29.86	**6.55**	124
	Cond-DBC	**27.65**	6.72	52
Medium	BC	27.80	6.01	11
	DBC	15.01	4.51	**7**
	Cond-DBC	**13.97**	**4.01**	9
Large	BC	17.12	5.20	7
	DBC	**9.59**	**2.87**	10
	Cond-DBC	12.37	3.73	**0**

Notably, at the small scale (117 K params), Cond-DBC already outperforms the medium BC baseline on L-shape FID (23.87 vs. 31.25) and approaches the large BC model (20.45), indicating that diffusion guidance substantially compensates for limited network capacity. These results have direct practical implications: a 16× parameter reduction from 31 M to 1.9 M enables deployment on resource-constrained edge devices common in warehouse AMR fleets, while achieving superior path fidelity through a single forward pass at inference time.

5 CONCLUSION

SPADE bridges the generalization–efficiency gap: a medium Cond-DBC model (1.9 M params) achieves 39.1% lower APE and 33.5% lower FID than the large BC baseline (31 M params), while preserving real-time, on-edge inference. The key enabler is conditional diffusion guidance that models $p(a|s)$ directly, providing environment-aware training signals that benefit even capacity-limited networks. While our evaluation focuses on two expert-validated trajectory families, the geometry-agnostic formulation readily extends to additional shapes. Future work targets dynamic obstacle handling and annotation tool user studies.

REFERENCES

[1] Joao Carvalho, An T. Le, Mark Baierl, Dorothea Koert, and Jan Peters. 2024. Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models. arXiv:2308.01557 [cs.RO] https://arxiv.org/abs/2308.01557

[2] Shang-Fu Chen, Hsiang-Chun Wang, Ming-Hao Hsu, Chun-Mao Lai, and Shao-Hua Sun. 2024. Diffusion Model-Augmented Behavioral Cloning. arXiv:2302.13335 [cs.LG] https://arxiv.org/abs/2302.13335

[3] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. 2024. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. arXiv:2303.04137 [cs.RO] https://arxiv.org/abs/2303.04137

[4] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics* 4, 2 (1968), 100–107. https://doi.org/10.1109/TSSC.1968.300136

[5] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. arXiv:2006.11239 [cs.LG] https://arxiv.org/abs/2006.11239

[6] S.M. LaValle and J.J. Kuffner. 1999. Randomized kinodynamic planning. In *Proceedings 1999 IEEE International Conference on Robotics and Automation.* 473–479. https://doi.org/10.1109/ROBOT.1999.770022

[7] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. 2017. FiLM: Visual Reasoning with a General Conditioning Layer. arXiv:1709.07871 [cs.CV] https://arxiv.org/abs/1709.07871

[8] Mark Pfeiffer, Michael Schaeuble, Juan Nieto, Roland Siegwart, and Cesar Cadena. 2017. From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. https://doi.org/10.1109/icra.2017.7989182

[9] Anthony Rizk, Anthony Rizk, Charbel Abi Hana, Youssef Bakouny, and Flavia Khatounian. 2025. End-to-end Sketch-Guided Path Planning through Imitation Learning for Autonomous Mobile Robots. In *2025 IEEE 6th International Conference on Image Processing, Applications and Systems (IPAS)*. 1–7. https://doi.org/10.1109/IPAS63548.2025.10924509

[10] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. arXiv:1505.04597 [cs.CV] https://arxiv.org/abs/1505.04597

[11] Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. 2011. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. arXiv:1011.0686 [cs.LG] https://arxiv.org/abs/1011.0686

[12] David Silver, J. Andrew Bagnell, and Anthony Stentz. 2010. Learning from Demonstration for Autonomous Navigation in Complex Unstructured Terrain. *The International Journal of Robotics Research* 29, 12 (2010), 1565–1592. https://doi.org/10.1177/0278364910369715

[13] Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. Behavioral Cloning from Observation. arXiv:1805.01954 [cs.AI] https://arxiv.org/abs/1805.01954

[14] Xuesu Xiao, Bo Liu, Garrett Warnell, Jonathan Fink, and Peter Stone. 2020. APPLD: Adaptive Planner Parameter Learning From Demonstration. *IEEE Robotics and Automation Letters* 5, 3 (2020), 4541–4547. https://doi.org/10.1109/LRA.2020.3002217

[15] Maryam Zare, Parham M. Kebria, Abbas Khosravi, and Saeid Nahavandi. 2024. A Survey of Imitation Learning: Algorithms, Recent Developments, and Challenges. *IEEE Transactions on Cybernetics* 54, 12 (2024), 7173–7186. https://doi.org/10.1109/TCYB.2024.3395626