

StereoPolicy: Improving Robotic Manipulation Policies via Stereo Perception

Evans Han^{1,2} Yunfan Jiang¹ Yingke Wang^{1,*} Haoyue Xiao^{1,*}
 Huang Huang¹ Jianwen Xie³ Jiajun Wu¹ Li Fei-Fei¹ Ruohan Zhang^{1,2}
¹Stanford University ²Northwestern University ³Lambda, Inc

Abstract: Recent advances in robot imitation learning have produced powerful visuomotor policies that manipulate diverse objects from visual inputs. However, monocular observations lack depth information, which is critical for precise manipulation in cluttered or geometrically complex scenes. Explicit depth maps and point clouds are often noisy and fragile in real-world manipulation. We introduce **StereoPolicy**, a visuomotor policy learning framework that directly leverages synchronized stereo image pairs to improve geometric reasoning without constructing explicit 3D representations. StereoPolicy processes each image with pre-trained 2D vision encoders and fuses left-right features through a cross-attention-based Stereo Transformer, capturing spatial correspondence and disparity cues implicitly. The framework integrates with diffusion-based and pretrained vision-language-action (VLA) policies, delivering consistent improvements over RGB, RGB-D, point cloud, and multi-view baselines across three simulation benchmarks and seven real-robot tabletop and bimanual mobile manipulation tasks. Our results show that stereo vision bridges 2D pretrained representations and 3D geometric understanding for robotic manipulation. The project page and additional details are available at <https://stereopolicy.github.io>.

Keywords: Robotic Perception, Stereo Vision, Robot Manipulation




Figure 1: Compared to traditional visual modalities for robot learning, stereo input provides certain advantages. This work investigates how we best leverage stereo input for policy learning.

*Denotes equal contribution. Correspondence to Evans Han <evanshan@stanford.edu>.

1 Introduction

Large-scale robot learning has recently produced impressive visuomotor policies and vision-language-action (VLA) models that perform diverse manipulation tasks from visual inputs alone [1–12]. These policies typically operate on monocular images processed by pretrained 2D vision encoders to extract semantic and spatial features for downstream action generation [13–15]. While monocular inputs can be effective in controlled scenes, they inherently lack accurate depth and geometric awareness. These capabilities are crucial for robots to operate in complicated scenes, such as daily environments [16–20], and to perform dexterous and fine-grained manipulation skills [21, 22].

Stereo vision instead naturally encodes 3D spatial structure through synchronized image pairs as shown in Figure 1. A long line of research [23–33] in computer vision has demonstrated that stereo pairs significantly improve depth estimation and scene geometry understanding compared to monocular images, suggesting strong potential for robotic applications where 3D information is necessary.

However, prior work in robotics typically employs stereo images indirectly, by first constructing explicit 3D representations such as depth maps [16, 17, 34, 35] and point clouds [18–21, 36–38], and then feeding these into control or planning modules. While effective in controlled settings, such approaches face practical challenges. First, explicit 3D representations are sensitive to sensor noise, calibration errors, and object properties, making high-quality reconstruction difficult to achieve reliably. Even with careful calibration, depth maps often miss fine semantic details, and point clouds suffer from object deformation, as visualized in Figure 1. Second, pretrained 3D vision encoders remain less mature than their 2D counterparts, largely due to the scarcity of large-scale 3D datasets [39–41] and scalable 3D backbone architectures [42–48]. Third, robotic manipulation requires both high precision and low latency, whereas stereo pipelines that produce explicit disparity maps [33] incur substantial inference overhead, making them ill-suited for real-time manipulation.

In this work, we take a different approach. Instead of relying on explicit 3D reconstruction, we design StereoPolicy, a framework that directly consumes synchronized stereo image pairs for action generation. Our method leverages 2D vision encoders pretrained on internet-scale data to extract features from each image independently, then fuses them into stereo-aware tokens through cross attention in a Stereo Transformer. This enables the policy to implicitly capture 3D spatial correspondence while retaining the generalization of pretrained 2D backbones, yielding stereo-aware visuomotor policies without explicit 3D supervision or reconstruction.

Our main results can be summarized as follows:

Performance. We empirically demonstrate that StereoPolicy outperforms policies based on RGB, RGB-D, point cloud, and multi-view across diverse real and simulated manipulation tasks, spanning tabletop and bimanual mobile settings. It can also be integrated to improve pretrained VLA models.

Hardware Insights. We systematically analyze key hardware factors including camera angles and stereo baseline-to-distance ratio, deriving practical guidelines for configuring stereo setups.

Model Design. We comprehensively study stereo encoder design, comparing different backbone choices and identifying an effective architecture for robust stereo-based manipulation.

2 Method

We propose StereoPolicy, a stereo perception module to extract geometry-aware features for robot policies (§2.1). The resulting representations can be seamlessly integrated into both diffusion policies and VLA models without modifying their backbone architectures. Below, we describe two models that form the conditions of our experiments (detailed in §3).

- STEREOPOLICY-DP integrates stereo encoder into diffusion policy, trained from scratch on each benchmark task (§2.2).
- STEREOPOLICY-VLA combines the stereo encoder with pre-trained VLA model and fine-tunes the system (§2.3).




Figure 2: **StereoPolicy Pipeline.** Stereo inputs are encoded by a vision backbone, fused with a Stereo Transformer, and applied to both diffusion-policy training and finetuning VLA baselines.

2.1 Stereo Feature Extraction

Stereo vision has shown superior performance in 3D understanding, including depth estimation [33] and robot grasping [49] in simulation environments. Building upon this observation, we design a StereoPolicy that leverages stereo image pairs for robot manipulation. Formally, the policy consists of a stereo vision encoder and a policy backbone. The stereo vision encoder takes N pairs of stereo images: $\mathcal{I}_{t,i} = (I_{t,i}^L, I_{t,i}^R)$, where (I_i^L, I_i^R) denote the left and right images captured by a stereo camera setup from i^{th} view at timestep t .

Stereo image encoding Each image is processed by a 2D vision encoder $V_i(\cdot)$ to produce single-view feature maps: $f_{t,i}^L, f_{t,i}^R$, $f \in \mathbb{R}^{H \times W \times C}$. Within each stereo pair, the left and right encoders share weights to preserve geometric consistency. To enrich geometric reasoning, we augment the image features with DINOv2 [50] embeddings, which provide strong monocular priors to disambiguate regions where epipolar correspondence is unreliable [33]:

$$\tilde{f}_{t,i} = \phi([f_{t,i} \| f_{t,i}^{\text{DINO}}]), \quad (1)$$

where $\phi(\cdot)$ is a linear projection aligning feature dimensions. Empirically, we observe that DINOv2 improves external views but degrades wrist-view performance, likely due to domain mismatch with egocentric imagery. Therefore, DINOv2 is applied only to external views. We ablate the architecture without DINOv2 in §3.3.

Stereo feature fusion To lift 2D image features with implicit depth cues, we adopt an alternating attention mechanism in Stereo Transformer: self-attention computes attention between pixels, while cross attention computes attention between left and right images. We apply 2D RoPE [51] to the query and key projections of cross attention to encourage explicit cross-view correspondence learning and spatial position reasoning [49].

$$\forall i \in \{1, \dots, N\} : \quad z_{t,i} = \text{StereoTrans}_i(\tilde{f}_{t,i}^L, \tilde{f}_{t,i}^R) \in \mathbb{R}^{d_z}, \quad (2)$$

Multi-view aggregation Stereo fusion is performed independently per camera view. The resulting embeddings are then concatenated (denoted as $\parallel$) with low-dimensional proprioceptive features s_t as the input to the policy backbone to predict robot actions:

$$z_{\text{obs}} = [z_{t,1} \| z_{t,2} \| \dots \| z_{t,N} \| s_t]. \quad (3)$$




Figure 3: **Real-World Task Visualizations.** *Top*: Five tabletop tasks, with the final state shown for each. Tasks 3–5 vary by cup texture. *Bottom*: Two mobile manipulation tasks.

2.2 STEREOPOLICY-DP: Diffusion Policy with StereoPolicy

For imitation learning, we primarily adopt a diffusion policy [8] as the policy backbone.

The policy predicts a sequence of continuous actions $\mathbf{a}_{t:t+H-1} = [\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+H-1}]$ over a planning horizon H based on z_{obs} . In STEREOPOLICY-DP, we inject stereo perception by conditioning the denoising network on the fused stereo embedding z_{obs} , which encodes cross-view geometric correspondence. Concretely, at each denoising step τ , the policy predicts noise added to action:

$$\hat{\epsilon} = \epsilon_{\theta}(x^{(\tau)}, \tau, z_{\text{obs}}), \quad (4)$$

The model is trained with standard denoising objective and performs iterative refinement from Gaussian noise at inference time. This design allows STEREOPOLICY-DP to enhance action generation with implicit stereo cues while preserving the original diffusion policy architecture and training objective.

2.3 STEREOPOLICY-VLA: Adapting Monocular VLA to Stereo Inputs

Pre-trained VLA models exhibit strong semantic understanding through VLM pretraining, but their depth reasoning is limited by monocular inputs. To improve spatial perception, we extend the visual input from monocular to stereo by introducing a lightweight stereo feature extractor on top of the pre-trained vision encoder.

Concretely, we replace the monocular visual embedding with a stereo-aware representation. For each stereo image pair, we extract features with the pretrained vision encoder from VLA, then inject fused stereo representation from Stereo Transformer into the VLA backbone to predict actions:

$$\mathbf{a} = \pi_{\theta}(z_{\text{obs}}, x_{\text{lang}}, s), \quad (5)$$

where π_{θ} denotes the VLA policy (e.g., DiT), z_{obs} is fused stereo feature from StereoTrans, x_{lang} is the language instruction, and s denotes additional state.

3 Experiments

Our experiments were guided by the following research questions:




Figure 4: **Simulation Task Visualization**, from three benchmarks: OMNIGIBSON (4 tasks), ROBOCASA (24 tasks), and ROBOMIMIC (3 tasks).

- Q1. How does StereoPolicy perform compared to monocular RGB, RGB-D, point cloud, and multi-view-based policies?
- Q2. Can StereoPolicy be readily combined with large pre-trained policy models, such as vision-language-action (VLA) models, although these models are trained on monocular RGB data?
- Q3. How do factors such as stereo baseline (distance between two cameras), distance to the objects, and camera angles affect StereoPolicy’s performance?
- Q4. Which design choices, such as vision backbone choice and attention mechanism design, affect the best performance in StereoPolicy?

3.1 Tasks and Environments

We evaluate StereoPolicy on two robotic embodiments: a 7-DOF tabletop Franka arm and a 14-DOF bimanual mobile manipulation robot Galaxea R1 Pro, in both real-world and simulated settings. Task and data collection details are summarized in [Appendix C](#).

Real-World We design seven tasks spanning both tabletop and mobile manipulation ([Figure 3](#)).

Simulation We include tasks from three widely-used robotics benchmarks, including ROBOMIMIC [52], ROBOCASA [53], and four customized tasks in OMNIGIBSON [54]. [Figure 4](#) visualizes tasks.

3.2 Implementation Details

The StereoPolicy architecture, baseline implementations, training details, and evaluation protocol are described in [Appendix B](#).

Baselines We compare STEREOPOLICY-DP against six baseline architectures: monocular RGB, multi-view RGB, RGB-D, RGBD-3DDA [55], point-clouds with PointNet [56], and point-clouds with DP3 [37]. All baselines use the same policy backbones and training budgets where applicable.

To effectively evaluate STEREOPOLICY-VLA, we adopt 1) **Pi0.5** [57] and 2) **GR00T N1.5** [58] as our baseline VLA frameworks and follow the same settings from the original implementation.

3.3 Experiment Results

(Q1) **STEREOPOLICY-DP consistently outperforms monocular RGB, RGB-D, point cloud, and multi-view-based policies.** [Table 1](#) and [Table 2](#) compare STEREOPOLICY-DP with standard visual modalities in real-robot and simulation respectively. Overall, STEREOPOLICY-DP presents strong performance in both settings. In particular, STEREOPOLICY-DP outperforms RGB across all tasks, highlighting the importance of geometric reasoning for manipulation. For less depth-sensitive tasks (e.g., *Transport*), monocular RGB inputs already provide sufficient visual information for

Method	BANANA PNP	TOAST INSERT	PLASTIC CUP HANG	STEEL CUP HANG	GLASS CUP HANG	AVG SR (%)
RGB	12/20	7/20	12/20	10/20	1/20	42.0%
RGB-D	14/20	8/20	11/20	8/20	0/20	41.0%
RGBD-3DDA	13/20	9/20	13/20	10/20	0/20	45.0%
PCD-PointNet	7/20	0/20	5/20	2/20	0/20	14.0%
PCD-DP3	11/20	3/20	8/20	5/20	0/20	27.0%
Multi-View	13/20	8/20	13/20	9/20	1/20	44.0%
STEREOPOLICY-DP	16/20	12/20	15/20	13/20	3/20	59.0%

Table 1: **Real-World Tabletop Task Performance.** STEREOPOLICY-DP consistently outperforms other visual modalities. PCD performs worst in all real tasks; both RGB-D and PCD fail on glass cup hang tasks, demonstrating that explicit 3D methods are sensitive to object texture in practice.

Benchmark	OMNIGIBSON								ROBOMIMIC					
Task	STRAWBERRY		POURWATER		OPENDOOR		TURNONRADIO		TOOLHANG		SQUARE		TRANSPORT	
# Demos	100	200	100	200	200	300	200	300	100	200	100	200	100	200
RGB	59 _{3.4}	88 _{0.0}	10 _{4.3}	46 _{2.8}	26 _{2.8}	77 _{1.9}	42 _{1.6}	71 _{2.5}	53 _{1.9}	90 _{0.9}	74 _{1.6}	98 _{0.0}	92 _{1.6}	94 _{2.8}
RGB-D	63 _{3.8}	85 _{4.7}	16 _{1.6}	52 _{0.0}	31 _{2.5}	80 _{4.3}	47 _{0.9}	73 _{3.4}	56 _{0.0}	88 _{2.5}	79 _{1.9}	92 _{0.0}	94 _{1.6}	94 _{1.6}
RGBD-3DDA	74 _{4.3}	93 _{1.9}	26 _{0.0}	61 _{3.4}	48 _{3.3}	100 _{0.0}	46 _{4.3}	75 _{4.1}	84 _{1.6}	92 _{2.8}	83 _{1.9}	97 _{2.5}	94 _{1.6}	96 _{1.6}
PCD-DP3	45 _{3.4}	63 _{2.5}	3 _{4.7}	31 _{0.9}	30 _{1.6}	69 _{4.1}	35 _{1.9}	64 _{4.9}	40 _{1.6}	76 _{0.9}	69 _{2.5}	88 _{1.6}	63 _{2.5}	72 _{2.8}
Multi-View	68 _{3.3}	89 _{2.5}	21 _{0.9}	52 _{3.3}	31 _{4.7}	75 _{3.8}	43 _{3.4}	71 _{1.9}	54 _{0.0}	92 _{0.9}	78 _{0.0}	96 _{1.6}	92 _{2.8}	94 _{0.0}
STEREOPOLICY-DP	82_{2.8}	100_{0.0}	34_{1.6}	70_{2.5}	57_{1.9}	100_{0.0}	55_{3.4}	82_{4.3}	94_{1.9}	96_{2.8}	88_{0.9}	100_{0.0}	94_{0.0}	96_{1.6}

Table 2: **Simulation Task Performance of Diffusion Policies across Visual Modalities.** Stereo input consistently improves performance, especially under low-data regime. We report average scores and standard deviations across three seed runs. All numbers are reported in percentage (%).

successful simulation. For tasks involving occlusion and fine-grained 3D alignment (e.g., *ToolHang* and *PourWater*), STEREOPOLICY-DP shows clear improvements ¹.

However, point cloud-based methods perform poorly overall: real-world depth measurements are often noisy (See Figure 7), and even in simulation DP3 still underperforms on most benchmarks. We observe point-cloud methods are fragile and heavily rely on careful pre-processing. Without careful cropping and sampling, its performance can easily drop to 0. While geometry-aware models such as RGBD-3DDA achieve competitive results on simulation due to perfect depth, their performance degrades notably on reflective and transparent objects (e.g., steel and glass cup hang), where depth sensing becomes unreliable; its reliance on calibrated depth sensors also limits robustness. In contrast, STEREOPOLICY-DP learns generalizable geometry directly from image pairs and stays robust in challenging settings. STEREOPOLICY-DP also surpasses 3DDA in simulation, indicating that our stereo architecture can not only recover geometric information on par with explicit depth sensors, but also leverage implicit stereo cues to improve fine-grained control.




Figure 7: RGB-D and PCD are fragile in real environment. Glass cup is entirely missing.

depth, their performance degrades notably on reflective and transparent objects (e.g., steel and glass cup hang), where depth sensing becomes unreliable; its reliance on calibrated depth sensors also limits robustness. In contrast, STEREOPOLICY-DP learns generalizable geometry directly from image pairs and stays robust in challenging settings. STEREOPOLICY-DP also surpasses 3DDA in simulation, indicating that our stereo architecture can not only recover geometric information on par with explicit depth sensors, but also leverage implicit stereo cues to improve fine-grained control.

STEREOPOLICY-DP outperforms multi-view-based policies Multi-view based policies also mounts stereo cameras, but does not pass through an additional Stereo Transformer module, validating the effectiveness of our stereo architecture.

STEREOPOLICY-DP is more data-efficient. In OMNIGIBSON tasks, STEREOPOLICY-DP reaches higher SR with fewer demonstrations. Although Peri et al. [59] claims point-cloud inputs can yield higher sample efficiency compared to RGB/RGB-D, here we observe the opposite.

(Q2) StereoPolicy enhances pretrained VLA models, although these models are trained on monocular data. We next examine the STEREOPOLICY-VLA, where StereoPolicy is incorporated into a state-of-the-art pre-trained VLA model, Pi0.5 and GR00T-N1.5 for fine-tuning. As

¹We present the inference latency of RGB, RGBD-3DDA and STEREOPOLICY-DP in Appendix E.




Figure 5: Average STEREOPOLICY-VLA Performance on ROBOCASA-KITCHEN 24 Tasks.






Figure 6: STEREOPOLICY-VLA (Pi0.5) Performance on bimanual mobile manipulation tasks in both real-world and simulation.

summarized in Figure 5, incorporating stereo input and stereo fusion module yields consistent improvements across both high-data (300 and 100 demos) and low-data (30 demos) settings. These gains, while modest, confirm our earlier claim: stereo features provide geometric cues that pre-trained VLM encoders inherently lack, and fine-tuning allows these large models to benefit from stereo depth information without disrupting their learned semantic representations.

Figure 6 further evaluates STEREOPOLICY-VLA on mobile manipulation. STEREOPOLICY-VLA outperforms monocular baselines in both simulation and real-world settings. In particular, RGB policies fail on precise actions such as inserting the gripper into a radio handle and pressing buttons accurately (failure case visualization in Appendix), while STEREOPOLICY-VLA succeeds.



(a) **Camera Angle Visualization.** From left to right: *sideview left, agentview left, frontview, agentview right, and sideview right.*



(b) **SR across 10 Different Camera Angle Views** over TOOLHANG task for RGB policy and STEREOPOLICY-DP. STEREOPOLICY-DP is more robust towards changing camera angles than RGB baseline. Appendix B provides full mapping of camera indices.

Figure 8: Performance of STEREOPOLICY-DP across different camera angles.

(Q3) STEREOPOLICY-DP is most effective when the baseline is approximately 10% of the target object distance. We vary the stereo baseline distance (2cm, 6cm, 10cm) and camera-object distance (0.6m–1.0m) while keeping other conditions fixed. Results show that performance is governed not by either factor alone, but by their ratio:

$$r = \frac{\text{stereo baseline}}{\text{camera-object distance}} \quad (6)$$

which determines the effective disparity. STEREOPOLICY-DP achieves the best performance when r lies in the range of $r \in [0.09, 0.13]$, as shown in Figure 9. Within this range, disparity is strong enough for reliable depth estimation while preserving sufficient overlap between left and right views. For typical tabletop setups (0.6–0.8m camera-object distance), a 6 cm baseline performs best.

When r becomes too small (e.g., 2 cm baseline or > 0.9 m distance), performance drops due to weak disparity and poor depth discrimination. Conversely, excessively large r values (e.g., 10 cm baseline at 0.6 m) reduce view overlap and cause geometric inconsistency across views. In conclusion, stereo policies are most effective when the baseline is approximately 10% of the target distance. This ratio maintains an optimal, balanced trade-off between disparity and overlap. It allows STEREOPOLICY-DP to extract stable geometric features and improves manipulation accuracy across viewpoints.




Figure 9: Effect of Stereo Baseline and Distance on Task Success Rate. Figure 10: Vision Encoder and Component Ablation. Experiments on ToolHang task with 100 demos.

In addition, STEREOPOLICY-DP consistently outperforms RGB-based policies across different camera angles. We evaluate 10 camera angles on the TOOLHANG task, visualized in Figure 8a; Figure 8b shows the corresponding views and performance. The largest gains appear in front-facing views (*frontview*, +18%; *frontview-high*, +14%), where stereo disparity provides complementary depth cues. Oblique views such as *agentview-left/right* and their elevated counterparts also benefit, as stereo helps resolve partial occlusions and spatial correspondence. Side views show smaller gains due to limited overlap and stronger self-occlusion. Overall, stereo is most helpful when monocular cues are geometrically ambiguous.

(Q4) The choice of vision backbone significantly influences STEREOPOLICY-DP’s performance, particularly in low-data regimes. To explore the most effective vision backbone for robot manipulation, we evaluate several architectures on the TOOLHANG task. As shown in Figure 10, in low data regime (under 100 demonstrations), larger pretrained vision-language backbones (e.g., OPENCLIP-B/16 and SIGLIP-SO400M/14) yield substantial performance gains over smaller convolutional architectures such as RESNET18. These models provide richer representations and generalize better with limited supervision. Multi-scale features from architecture like EdgeNeXt-S [60] or Resnet18 with Feature Pyramid Network [61] also significantly benefit Stereo Transformer, indicating that feature hierarchy is important for correspondence learning. From the ablation, we observe that concatenating external view image features from frozen DINOv2 further improves performance. Empirically, we observe the gains are view-dependent: it improves external-view but hurts wrist-view performance, likely due to domain mismatch. Removing the Stereo Transformer drops performance from 0.94 to 0.85, validating the effectiveness of stereo fusion module.

4 Conclusion

We presented StereoPolicy, a visuomotor learning framework that directly consumes synchronized stereo image pairs instead of explicit 3D reconstructions. By combining pretrained 2D vision encoders with a cross-attention-based stereo fusion module, StereoPolicy improves over monocular RGB, RGB-D, point-cloud, and multi-view baselines across both simulated and real-world settings, spanning tabletop and mobile manipulation tasks. Our analysis further identifies practical design choices for stereo-based policies, including camera baseline, viewpoint, vision backbone, and fusion architecture. StereoPolicy also improves pretrained VLA models, suggesting that stereo cues can complement large-scale 2D visual representations for robot learning.

5 Limitations and Future Work

While promising, our study is limited to relatively small-scale tabletop and mobile manipulation settings. Although StereoPolicy improves over monocular RGB, RGB-D, point-cloud, and multi-view baselines on transparent and reflective objects, these scenarios remain challenging in absolute terms, and the overall success rate is still relatively low. We observe that performance can be sensitive to lighting, reflections, and stereo camera placement. Future work could scale StereoPolicy to larger robot datasets with stereo pairs (e.g. DROID [62]) and extend it to more embodiments.

References

- [1] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies, 2016. URL <https://arxiv.org/abs/1504.00702>.
- [2] Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, and N. Heess. Reinforcement and imitation learning for diverse visuomotor skills, 2018. URL <https://arxiv.org/abs/1802.09564>.
- [3] M. Shridhar, L. Manuelli, and D. Fox. Cliport: What and where pathways for robotic manipulation. *arXiv preprint arXiv: Arxiv-2109.12098*, 2021.
- [4] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, T. Eccles, J. Bruce, A. Razavi, A. Edwards, N. Heess, Y. Chen, R. Hadsell, O. Vinyals, M. Bordbar, and N. de Freitas. A generalist agent. *arXiv preprint arXiv: Arxiv-2205.06175*, 2022.
- [5] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan. Vima: General robot manipulation with multimodal prompts. *arXiv preprint arXiv: 2210.03094*, 2022.
- [6] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. C. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. R. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. Rt-1: Robotics transformer for real-world control at scale. *Robotics: Science and Systems*, 2022. doi:10.48550/arXiv.2212.06817. URL <https://arxiv.org/abs/2212.06817v2>.
- [7] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, K. Choromanski, T. Ding, D. Driess, K. A. Dubey, C. Finn, P. R. Florence, C. Fu, M. G. Arenas, K. Gopalakrishnan, K. Han, K. Hausman, A. Herzog, J. Hsu, B. Ichter, A. Irpan, N. J. Joshi, R. C. Julian, D. Kalashnikov, Y. Kuang, I. Leal, S. Levine, H. Michalewski, I. Mordatch, K. Pertsch, K. Rao, K. Reymann, M. Ryoo, G. Salazar, P. R. Sanketi, P. Sermanet, J. Singh, A. Singh, R. Soricut, H. Tran, V. Vanhoucke, Q. Vuong, A. Wahid, S. Welker, P. Wohlhart, T. Xiao, T. Yu, and B. Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *Conference on Robot Learning*, 2023. doi:10.48550/arXiv.2307.15818. URL <https://arxiv.org/abs/2307.15818v1>.
- [8] 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.
- [9] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In K. E. Bekris, K. Hauser, S. L. Herbert, and J. Yu, editors, *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. doi:10.15607/RSS.2023.XIX.016. URL <https://doi.org/10.15607/RSS.2023.XIX.016>.
- [10] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. OpenVLA: An open-source vision-language-action model. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=ZMnD6QZAE6>.
- [11] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv: 2410.24164*, 2024.

- [12] NVIDIA, J. Bjorck, F. Castañeda, N. Cherniadev, X. Da, R. Ding, L. J. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, J. Jang, Z. Jiang, J. Kautz, K. Kundalia, L. Lao, Z. Li, Z. Lin, K. Lin, G. Liu, E. Llontop, L. Magne, A. Mandlekar, A. Narayan, S. Nasiriany, S. Reed, Y. L. Tan, G. Wang, Z. Wang, J. Wang, Q. Wang, J. Xiang, Y. Xie, Y. Xu, Z. Xu, S. Ye, Z. Yu, A. Zhang, H. Zhang, Y. Zhao, R. Zheng, and Y. Zhu. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv: 2503.14734*, 2025.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition, 2015. URL <https://arxiv.org/abs/1512.03385>.
- [14] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [15] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. URL <https://arxiv.org/abs/2010.11929>.
- [16] M. Liu, Z. Chen, X. Cheng, Y. Ji, R.-Z. Qiu, R. Yang, and X. Wang. Visual whole-body control for legged loco-manipulation. *arXiv preprint arXiv: 2403.16967*, 2024.
- [17] S. Uppal, A. Agarwal, H. Xiong, K. Shaw, and D. Pathak. Spin: Simultaneous perception, interaction and navigation. *CVPR*, 2024.
- [18] J. Yang, Z. ang Cao, C. Deng, R. Antonova, S. Song, and J. Bohg. Equibot: Sim(3)-equivariant diffusion policy for generalizable and data efficient learning, 2024. URL <https://arxiv.org/abs/2407.01479>.
- [19] Y. Jiang, R. Zhang, J. Wong, C. Wang, Y. Ze, H. Yin, C. Gokmen, S. Song, J. Wu, and L. Fei-Fei. BEHAVIOR robot suite: Streamlining real-world whole-body manipulation for everyday household activities. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=v2KevjWScT>.
- [20] P. Sundaresan, R. Malhotra, P. Miao, J. Yang, J. Wu, H. Hu, R. Antonova, F. Engelmann, D. Sadigh, and J. Bohg. Homer: Learning in-the-wild mobile manipulation via hybrid imitation and whole-body control, 2025. URL <https://arxiv.org/abs/2506.01185>.
- [21] Y. Qin, B. Huang, Z.-H. Yin, H. Su, and X. Wang. Dexpoint: Generalizable point cloud reinforcement learning for sim-to-real dexterous manipulation. *Conference on Robot Learning*, 2022. doi:10.48550/arXiv.2211.09423.
- [22] C. Wang, H. Shi, W. Wang, R. Zhang, F.-F. Li, and K. Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *ROBOTICS*, 2024. doi:10.48550/arXiv.2403.07788.
- [23] D. Marr and T. Poggio. Cooperative computation of stereo disparity. *Science*, 194(4262):283–287, 1976. doi:10.1126/science.968482. URL <https://www.science.org/doi/abs/10.1126/science.968482>.
- [24] F. Zhang, X. Qi, R. Yang, V. Prisacariu, B. Wah, and P. Torr. Domain-invariant stereo matching networks. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 420–439, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58536-5.
- [25] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

- [26] H. Xu and J. Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1956–1965, 2020.
- [27] L. Lipson, Z. Teed, and J. Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *2021 International Conference on 3D Vision (3DV)*, pages 218–227, 2021. doi: [10.1109/3DV53792.2021.00032](https://doi.org/10.1109/3DV53792.2021.00032).
- [28] Z. Li, X. Liu, N. Drenkow, A. Ding, F. X. Creighton, R. H. Taylor, and M. Unberath. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6197–6206, October 2021.
- [29] Z. Shen, Y. Dai, and Z. Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13901–13910, 2021.
- [30] J. Li, P. Wang, P. Xiong, T. Cai, Z. Yan, L. Yang, J. Liu, H. Fan, and S. Liu. Practical stereo matching via cascaded recurrent network with adaptive correlation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16263–16272, June 2022.
- [31] P. Weinzaepfel, T. Lucas, V. Leroy, Y. Cabon, V. Arora, R. Brégier, G. Csurka, L. Antsfeld, B. Chidlovskii, and J. Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17969–17980, October 2023.
- [32] G. Xu, X. Wang, X. Ding, and X. Yang. Iterative geometry encoding volume for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21919–21928, June 2023.
- [33] B. Wen, M. Trepte, J. Aribido, J. Kautz, O. Gallo, and S. Birchfield. Foundationstereo: Zero-shot stereo matching, 2025. URL <https://arxiv.org/abs/2501.09898>.
- [34] K. Shankar, M. Tjersland, J. Ma, K. Stone, and M. Bajracharya. A learned stereo depth system for robotic manipulation in homes, 2021. URL <https://arxiv.org/abs/2109.11644>.
- [35] R. Yang, G. Yang, and X. Wang. Neural volumetric memory for visual locomotion control. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1430–1440, 2023.
- [36] A. Goyal, J. Xu, Y. Guo, V. Blukis, Y.-W. Chao, and D. Fox. Rvt: Robotic view transformer for 3d object manipulation. *ArXiv*, abs/2306.14896, 2023.
- [37] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. *arXiv preprint arXiv: 2403.03954*, 2024. URL <https://arxiv.org/abs/2403.03954v7>.
- [38] Y. Jiang, C. Wang, R. Zhang, J. Wu, and L. Fei-Fei. Transic: Sim-to-real policy transfer by learning from online correction. *arXiv preprint arXiv: 2405.10315*, 2024. URL <https://arxiv.org/abs/2405.10315v3>.
- [39] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. doi:[10.1109/CVPR.2012.6248074](https://doi.org/10.1109/CVPR.2012.6248074).
- [40] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

- [41] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In X. Jiang, J. Hornegger, and R. Koch, editors, *Pattern Recognition*, pages 31–42, Cham, 2014. Springer International Publishing. ISBN 978-3-319-11752-2.
- [42] C. Qi, H. Su, K. Mo, and L. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Computer Vision and Pattern Recognition*, 2016. doi:10.1109/CVPR.2017.16.
- [43] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [44] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [45] H. Zhao, L. Jiang, J. Jia, P. H. S. Torr, and V. Koltun. Point transformer. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16239–16248, 2020.
- [46] A. Jaegle, F. Gimeno, A. Brock, A. Zisserman, O. Vinyals, and J. Carreira. Perceiver: General perception with iterative attention. *arXiv preprint arXiv: 2103.03206*, 2021.
- [47] X. Wu, Y. Lao, L. Jiang, X. Liu, and H. Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 33330–33342. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/d78ece6613953f46501b958b7bb4582f-Paper-Conference.pdf.
- [48] X. Wu, L. Jiang, P.-S. Wang, Z. Liu, X. Liu, Y. Qiao, W. Ouyang, T. He, and H. Zhao. Point transformer v3: Simpler, faster, stronger. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4840–4851, 2023.
- [49] R. Singh, A. Allshire, A. Handa, N. Ratliff, and K. Van Wyk. Dextrah-rgb: Visuomotor policies to grasp anything with dexterous hands. *arXiv preprint arXiv:2412.01791*, 2024.
- [50] M. Oquab and et al. Dinov2: Learning robust visual features without supervision, 2024. URL <https://arxiv.org/abs/2304.07193>.
- [51] B. Heo, S. Park, D. Han, and S. Yun. Rotary position embedding for vision transformer, 2024. URL <https://arxiv.org/abs/2403.13298>.
- [52] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *arXiv preprint arXiv:2108.03298*, 2021.
- [53] S. Nasiriany, A. Maddukuri, L. Zhang, A. Parikh, A. Lo, A. Joshi, A. Mandlekar, and Y. Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. In *Robotics: Science and Systems (RSS)*, 2024.
- [54] C. Li, R. Zhang, J. Wong, C. Gokmen, S. Srivastava, R. Martín-Martín, C. Wang, G. Levine, W. Ai, B. Martinez, H. Yin, M. Lingelbach, M. Hwang, A. Hiranaka, S. Garlanka, A. Aydin, S. Lee, J. Sun, M. Anvari, M. Sharma, D. Bansal, S. Hunter, K.-Y. Kim, A. Lou, C. R. Matthews, I. Villa-Renteria, J. H. Tang, C. Tang, F. Xia, Y. Li, S. Savarese, H. Gweon, C. K. Liu, J. Wu, and L. Fei-Fei. Behavior-1k: A human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation. *arXiv preprint arXiv:2403.09227*, 2024.
- [55] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations, 2024. URL <https://arxiv.org/abs/2402.10885>.

- [56] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation, 2017. URL <https://arxiv.org/abs/1612.00593>.
- [57] P. Intelligence, K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, M. Y. Galliker, D. Ghosh, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, D. LeBlanc, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, A. Z. Ren, L. X. Shi, L. Smith, J. T. Springenberg, K. Stachowicz, J. Tanner, Q. Vuong, H. Walke, A. Walling, H. Wang, L. Yu, and U. Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization, 2025. URL <https://arxiv.org/abs/2504.16054>.
- [58] NVIDIA, N. C. Johan Bjorck and Fernando Castañeda, X. Da, R. Ding, L. J. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, J. Jang, Z. Jiang, J. Kautz, K. Kundalia, L. Lao, Z. Li, Z. Lin, K. Lin, G. Liu, E. Llontop, L. Magne, A. Mandlekar, A. Narayan, S. Nasiriany, S. Reed, Y. L. Tan, G. Wang, Z. Wang, J. Wang, Q. Wang, J. Xiang, Y. Xie, Y. Xu, Z. Xu, S. Ye, Z. Yu, A. Zhang, H. Zhang, Y. Zhao, R. Zheng, and Y. Zhu. GR00T N1: An open foundation model for generalist humanoid robots. In *ArXiv Preprint*, March 2025.
- [59] S. Peri, I. Lee, C. Kim, L. Fuxin, T. Hermans, and S. Lee. Point cloud models improve visual robustness in robotic learners, 2024. URL <https://arxiv.org/abs/2404.18926>.
- [60] M. Maaz, A. Shaker, H. Cholakkal, S. Khan, S. W. Zamir, R. M. Anwer, and F. S. Khan. Ed-genext: Efficiently amalgamated cnn-transformer architecture for mobile vision applications, 2022. URL <https://arxiv.org/abs/2206.10589>.
- [61] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection, 2017. URL <https://arxiv.org/abs/1612.03144>.
- [62] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [63] O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec, S. Yi, M. Ramamonjisoa, F. Massa, D. Haziza, L. Wehrstedt, J. Wang, T. Darcet, T. Moutakanni, L. Sentana, C. Roberts, A. Vedaldi, J. Tolan, J. Brandt, C. Couprie, J. Mairal, H. Jégou, P. Labatut, and P. Bojanowski. Dinov3, 2025. URL <https://arxiv.org/abs/2508.10104>.
- [64] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>.
- [65] X. Wu, D. DeTone, D. Frost, T. Shen, C. Xie, N. Yang, J. Engel, R. Newcombe, H. Zhao, and J. Straub. Sonata: Self-supervised learning of reliable point representations, 2025. URL <https://arxiv.org/abs/2503.16429>.
- [66] J. Hou, B. Graham, M. Nießner, and S. Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts, 2021. URL <https://arxiv.org/abs/2012.09165>.
- [67] X. Wu, X. Wen, X. Liu, and H. Zhao. Masked scene contrast: A scalable framework for unsupervised 3d representation learning, 2023. URL <https://arxiv.org/abs/2303.14191>.
- [68] X. Wu, L. Jiang, P.-S. Wang, Z. Liu, X. Liu, Y. Qiao, W. Ouyang, T. He, and H. Zhao. Point transformer v3: Simpler, faster, stronger, 2024. URL <https://arxiv.org/abs/2312.10035>.
- [69] S. Xie, J. Gu, D. Guo, C. R. Qi, L. J. Guibas, and O. Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding, 2020. URL <https://arxiv.org/abs/2007.10985>.

- [70] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud. Dust3r: Geometric 3d vision made easy, 2024. URL <https://arxiv.org/abs/2312.14132>.
- [71] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis, 2020. URL <https://arxiv.org/abs/2003.08934>.
- [72] Z. Fan, P. Wang, Y. Jiang, X. Gong, D. Xu, and Z. Wang. Nerf-sos: Any-view self-supervised object segmentation on complex scenes, 2022. URL <https://arxiv.org/abs/2209.08776>.
- [73] Y. Liu, L. Kong, J. Cen, R. Chen, W. Zhang, L. Pan, K. Chen, and Z. Liu. Segment any point cloud sequences by distilling vision foundation models, 2023. URL <https://arxiv.org/abs/2306.09347>.
- [74] A. Haque, M. Tancik, A. A. Efros, A. Holynski, and A. Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions, 2023. URL <https://arxiv.org/abs/2303.12789>.
- [75] K. Liu, F. Zhan, J. Zhang, M. Xu, Y. Yu, A. E. Saddik, C. Theobalt, E. Xing, and S. Lu. Weakly supervised 3d open-vocabulary segmentation, 2024. URL <https://arxiv.org/abs/2305.14093>.
- [76] J. Min, Y. Jeon, J. Kim, and M. Choi. S²M²: Scalable stereo matching model for reliable depth estimation, 2025. URL <https://arxiv.org/abs/2507.13229>.
- [77] L. Bartolomei, F. Tosi, M. Poggi, and S. Mattoccia. Stereo anywhere: Robust zero-shot deep stereo matching even where either stereo or mono fail, 2025. URL <https://arxiv.org/abs/2412.04472>.
- [78] G. Xu, X. Wang, Z. Zhang, J. Cheng, C. Liao, and X. Yang. Igev++: Iterative multi-range geometry encoding volumes for stereo matching, 2025. URL <https://arxiv.org/abs/2409.00638>.
- [79] L. Lipson, Z. Teed, and J. Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching, 2021. URL <https://arxiv.org/abs/2109.07547>.
- [80] J.-R. Chang and Y.-S. Chen. Pyramid stereo matching network, 2018. URL <https://arxiv.org/abs/1803.08669>.
- [81] Q. Wang, S. Shi, S. Zheng, K. Zhao, and X. Chu. Fadnet: A fast and accurate network for disparity estimation, 2020. URL <https://arxiv.org/abs/2003.10758>.
- [82] Q. Wang, S. Shi, S. Zheng, K. Zhao, and X. Chu. Fadnet++: Real-time and accurate disparity estimation with configurable networks, 2021. URL <https://arxiv.org/abs/2110.02582>.
- [83] Y. Wang, Y. Liang, Y. Hu, and Y. Fu. Robuststereo: Robust zero-shot stereo matching under adverse weather, 2025. URL <https://arxiv.org/abs/2507.01653>.
- [84] A. Kalra, V. Tamaazy, A. Dall’olio, R. Khanna, T. Gerlich, G. Giannopolou, G. Stoppi, D. Baxter, A. Ghosh, R. Szeliski, et al. A plentoptic 3d vision system. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–12, 2024.
- [85] V. Pradeep, C. Rhemann, S. Izadi, C. Zach, M. Bleyer, and S. Bathiche. Monofusion: Real-time 3d reconstruction of small scenes with a single web camera. In *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 83–88. IEEE, 2013.
- [86] K. Bai, H. Zeng, L. Zhang, Y. Liu, H. Xu, Z. Chen, and J. Zhang. Cleardepth: Enhanced stereo perception of transparent objects for robotic manipulation, 2025. URL <https://arxiv.org/abs/2409.08926>.

- [87] T. Kollar, M. Laskey, K. Stone, B. Thananjeyan, and M. Tjersland. Simnet: Enabling robust unknown object manipulation from pure synthetic data via stereo, 2021. URL <https://arxiv.org/abs/2106.16118>.
- [88] H. Li, T. Padir, and H. Jiang. Stereonavnet: Learning to navigate using stereo cameras with auxiliary occupancy voxels, 2024. URL <https://arxiv.org/abs/2403.12039>.
- [89] H. Li, Z. Li, N. U. Akmandor, H. Jiang, Y. Wang, and T. Padir. Stereovoxelnet: Real-time obstacle detection based on occupancy voxels from a stereo camera using deep neural networks, 2023. URL <https://arxiv.org/abs/2209.08459>.
- [90] T. G. W. Lum, M. Matak, V. Makoviychuk, A. Handa, A. Allshire, T. Hermans, N. D. Ratliff, and K. Van Wyk. Dextrah-g: Pixels-to-action dexterous arm-hand grasping with geometric fabrics. *arXiv preprint arXiv:2407.02274*, 2024.
- [91] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. URL <https://arxiv.org/abs/1505.04597>.
- [92] Y. Wu and K. He. Group normalization, 2018. URL <https://arxiv.org/abs/1803.08494>.
- [93] B. Sundaralingam, S. K. S. Hari, A. Fishman, C. Garrett, K. V. Wyk, V. Blukis, A. Millane, H. Oleynikova, A. Handa, F. Ramos, N. Ratliff, and D. Fox. curobo: Parallelized collision-free minimum-jerk robot motion generation, 2023.