


DynaMem: Online Dynamic Spatio-Semantic Memory for Open World Mobile Manipulation

Anonymous Author(s)

Affiliation

Address

email



1

Figure 1: An illustration of how our online dynamic spatio-semantic memory DynaMem responds to open vocabulary queries in a dynamic environment. During operation and exploration, DynaMem keeps updating its semantic map in memory. DynaMem maintains a voxelized pointcloud representation of the environment, and updates with dynamic changes in the environment by adding and removing points.

2

Abstract: Significant progress has been made in open-vocabulary mobile manipulation, where the goal is for a robot to perform tasks in any environment given a natural language description. However, most current systems assume a static environment, which limits the system’s applicability in real-world scenarios where environments frequently change due to human intervention or the robot’s own actions. In this work, we present DynaMem, a new approach to open-world mobile manipulation that uses a dynamic spatio-semantic memory to represent a robot’s environment. DynaMem constructs a 3D data structure to maintain a dynamic memory of point clouds, and answers open-vocabulary object localization queries using multimodal LLMs or open-vocabulary features generated by state-of-the-art vision-language models. Powered by DynaMem, our robots can explore novel environments, search for objects not found in memory, and continuously update the memory as objects move, appear, or disappear in the scene. We run extensive experiments on the Stretch SE3 robots in three real and nine offline scenes, and achieve an average pick-and-drop success rate of 70% on non-stationary objects, which is more than a $2\times$ improvement over state-of-the-art static systems.

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18 1 Introduction

19

Recent advances in robotics have made it possible to deploy robots in real world settings to tackle the open vocabulary mobile manipulation (OVMM) problem [1]. Here, the robots are tasked with navigating in unknown environments and interacting with objects following open vocabulary language instructions, such as “Pick up X from Y and put it in Z ”, where X , Y , and Z could be any object name or location. The two most common approaches to tackling OVMM are using policies trained in sim-

20

21

22

23

24 ulation and deploying them in the real world [2, 3, 4], or training modular systems that combine open
25 vocabulary navigation (OVN) [5, 6, 7, 8] with different robot manipulation skills [9, 10, 11, 12, 13].
26 Modular systems enjoy greater efficiency and success in real-world deployment [14] as they can di-
27 rectly leverage advances in vision and language models [9, 12], and are able to handle more diverse
28 and out-of-domain environments with no additional training.

29 However, as recent analysis has shown, the primary challenge in deploying modular OVMM is
30 that limitations of a module propagate to the entire system [9]. One key module in any OVMM
31 system is the open vocabulary navigation (OVN) module responsible for navigating to goals in the
32 environment. While many such OVN systems have been proposed in the literature [1, 8, 5, 11,
33 10, 9, 6, 7, 12, 13], they share a common limitation: they assume static, unchanging environments.
34 Contrast this with the real world, where environments change and objects are moved by either robots
35 or humans. Making such a restrictive assumption thus limits these systems’ applicability in real-
36 world settings. The primary reason behind this assumption is the lack of an effective dynamic
37 spatio-semantic memory that can adapt to both addition and removal of objects and obstacles in the
38 environment online.

39 In this work, we propose a novel spatio-semantic memory architecture, Dynamic 3D Voxel Memory
40 (DynaMem), that can adapt online to changes in the environment. DynaMem maintains a voxelized
41 pointcloud representation of an environment and adds or removes points as it observes the envi-
42 ronment change. Additionally, it supports two different ways to query the memory with natural
43 language: a vision-language model (VLM) featurized pointcloud, and a multimodal-LLM (mLLM)
44 QA system. Finally, DynaMem enables efficient exploration in changing environments by offering a
45 dynamic obstacle map and a value-based exploration map that can guide the robot to explore unseen,
46 outdated, or query-relevant parts of the world.

47 We evaluate DynaMem as a part of full open-vocabulary mobile manipulation stack in three real
48 world environments with multiple rounds of changes and manipulating multiple non-stationary ob-
49 jects, improving the static baseline by more than $2\times$ (70% vs. 30%). Additionally, we identify
50 an obstacle in efficiently developing dynamic spatio-semantic memory, namely the lack of dy-
51 namic benchmarks, since many OVN systems use static simulated environments [15, 16] or static
52 datasets [17, 18]. We address this by developing a new dynamic benchmark, DynaBench. It con-
53 sists of 9 different environments, each changing over time. We ablate our design choices in this
54 benchmark. To the best of our knowledge, DynaMem is the first spatio-semantic memory structure
55 supporting both adding and removing objects.

56 2 Method

57 In this section, we define our problem setup, and then describe our online, dynamic spatio-semantic
58 memory for open world, open vocabulary mobile manipulation. We introduce how to use this mem-
59 ory to localize text query and how to navigate to the target object in Appendix 6 and 7 respectively.

60 2.1 Problem Statement

61 We create our algorithm, DynaMem, to solve open vocabulary mobile manipulation (OVMM) prob-
62 lems in an open, constantly changing world. The goal in OVMM is for a mobile robot to execute a
63 series of manipulation commands given arbitrary language goals. We assume the following require-
64 ments for the memory module for dynamic, online operation:

- 65 • **Observations:** The mobile robot is equipped with an on-board RGB-D camera, and unlike prior
66 work [9], doesn’t start with a map of the environment. Rather, the robot explores the world and
67 use the online observed sequence of posed RGB-D images to build its map.
- 68 • **Environment dynamism:** The environment can change without the knowledge of the robot.

- 69 • **Localization queries:** Given a natural language query (i.e. "teddy bear"), the memory module
70 has to return the 3D location of the object or determine that the object doesn't exist in the scene
71 observed thus far.
- 72 • **Obstacle queries:** The memory module must determine whether a point in space is occupied by
73 an obstacle. Both the location of the objects and obstacles can move, previous observations often
74 contradict each other and must be resolved by the memory.

75 2.2 Dynamic 3D Voxel Map

76 Our answer to the challenge posed in the Section 2.1 is DynaMem. DynaMem is an evolving sparse
77 voxel map with associated information stored at each voxel, as shown in Fig. 6. In each non-empty
78 voxel, alongside its 3D location (x, y, z) , we also store the observation count C (how many times
79 that voxel was observed), source image ID I (which image the voxel was backprojected from), a
80 high-dimensional semantic feature vector f coming from a VLM like CLIP [19] or SigLIP [20], and
81 the latest observation time, t , in seconds.

82 To make this data structure dynamic, we describe the process with which we add and update with
83 new observations and remove outdated objects and associated voxels.

84 **Adding Points:** When the robot receives a new set of observations, i.e. RGB-D images with global
85 poses, we convert them to 3D coordinates in a global reference frame, and generate a semantic
86 feature vector for each point. The global coordinates are calculated from the global camera pose
87 and the backprojected depth image using the known camera transformation matrix. We calculate the
88 point-wise image feature by first converting the images to object patches by using a segmentation
89 model such as SAM-v2 [21], and then aggregating each patch feature over the output of a vision-
90 language models like CLIP [19] or SigLIP [20]. For more details about image-to-feature vector
91 mapping, we refer to earlier works [5, 9, 8]. Once we have calculated the points and associated
92 features, we cluster the new points and assign them to the nearest voxel grids. In Fig. 7, we show
93 how each voxel's metadata is updated. The count keeps track of the total number of assigned points
94 to each voxel grid, and the feature vector keeps track of the weighted average of all feature vectors
95 assigned to that voxel. Finally, the observation time and image ID are updated to keep track of the
96 latest observation contributing to a particular voxel. If a voxel was empty before assignment, we
97 assume its count $C = 0$ and feature vector $f = \vec{0}$.

98 **Removing Points:** When an object is moved or removed, its associated voxels in DynaMem may get
99 removed. We use ray-casting to find the outdated voxels. The operation follows a simple principle:
100 if a voxel falls within the frustum between the camera plane and the associated view point cloud,
101 that voxel must be unoccupied. To reduce the impact of the depth noise at long range, we don't
102 consider any pixel whose associated depth value is over 2m.

103 We illustrate a simplified 2D representation of this algorithm in Fig. 2. In practice, to speed up the
104 intersection between the sparse voxelmap and the view frustum, we project each existing voxel to
105 the camera plane and calculate the camera distance. If the image height and width are (H, W) , the
106 depth image is \mathbf{D} , and a certain voxel is projected to points (h, w) in the camera plane with depth d ,
107 it gets removed if both Eq. 1 and 2 hold.

$$(h, w) \in [0, H] \times [0, W] \tag{1}$$

$$d \in (0, \min(2, \mathbf{D}[h, w] + \epsilon)) \tag{2}$$

108 Where Eq. 1 ensures that the point falls within the camera view, and Eq 2 ensures that (a) the depth
109 $d > 0$, or the object is in front of camera, (b) $d < 2\text{m}$, or the voxel isn't too far away from the
110 camera, and (c) $d < \mathbf{D}[h, w]$ denoting the voxel is between the camera and the currently visible
111 object.

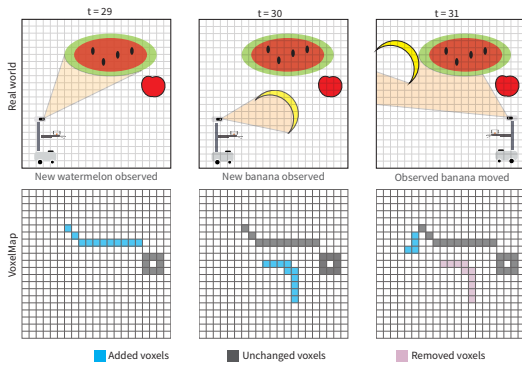


Figure 2: A high-level, 2D depiction of how adding and removing voxels from the voxel map works. New voxels are included which are in the RGB-D camera's view frustum, and old voxels that should block the view frustum but do not are removed from the map.



Figure 3: Querying DynaMem with a natural language query. First, we find the voxel with the highest alignment to the query. Next, we find the latest image of that voxel, and query with an open-vocabulary object detector to confirm the object location or abstain.

112 3 Experiments

113 We evaluate our method, DynaMem, on a Hello Robot: Stretch SE3 in real world environments. We
 114 also perform a series of ablation experiments in an offline benchmark in Appendix 8.

115 3.1 Real-world Experiments

116 As a baseline, we compare with OK-Robot [9], a state-of-the-art method for OVMM. OK-Robot
 117 uses a static voxelmap as its memory representation, and thus it highlights the importance of dy-
 118 namic memory for OVMM in a changing environment. For DynaMem, we run two variations of the
 119 algorithm in the real world: one with VLM-feature based queries and one with mLLM-QA based
 120 queries.

121 We describe detailed experiment setup in Appendix 9.

122 **Results:** Our experiments in three dynamic environments and with 30 queries is summarized in
 123 Fig. 4. We find that DynaMem with both VLM-feature based and mLLM-QA based queries have a
 124 total success rate of 70%. This is a significant improvement over the OK-Robot system, which has
 125 a total success rate of 30%. Notably, DynaMem is particularly adept at handling dynamic objects in
 126 the environment: only 6.7% of the trials failed due to our system not being able to navigate to such
 127 dynamic objects in the scene. This is in contrast to the OK-Robot system, where 53.3% of the trials
 128 failed because it could not find an object that moved in the environment. In contrast, navigating to
 129 static goals fails in only 10% of the cases for DynaMem with VLM-feature, 13.3% for OK-Robot
 130 and 20% for DynaMem with mLLM-QA.

131 4 Conclusions

132 In this work, we introduced DynaMem, a spatio-semantic memory for open-vocabulary mobile ma-
 133 nipulation that can handle changes to the environment during operation. We showed in three real
 134 world environments that DynaMem can navigate to, pick, and drop objects even while object and
 135 obstacle locations are changing.

References

- [1] S. Yenamandra, A. Ramachandran, K. Yadav, A. Wang, M. Khanna, T. Gervet, T.-Y. Yang, V. Jain, A. W. Clegg, J. Turner, Z. Kira, M. Savva, A. Chang, D. S. Chaplot, D. Batra, R. Mottaghi, Y. Bisk, and C. Paxton. Homerobot: Open-vocabulary mobile manipulation, 2024. URL <https://arxiv.org/abs/2306.11565>.
- [2] K. Ehsani, T. Gupta, R. Hendrix, J. Salvador, L. Weihs, K.-H. Zeng, K. P. Singh, Y. Kim, W. Han, A. Herrasti, et al. Imitating shortest paths in simulation enables effective navigation and manipulation in the real world. *arXiv preprint arXiv:2312.02976*, 2023.
- [3] R. Ramrakhya, D. Batra, E. Wijmans, and A. Das. Pirlnav: Pretraining with imitation and rl finetuning for objectnav. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17896–17906, June 2023.
- [4] K.-H. Zeng, Z. Zhang, K. Ehsani, R. Hendrix, J. Salvador, A. Herrasti, R. Girshick, A. Kembhavi, and L. Weihs. Poliformer: Scaling on-policy rl with transformers results in masterful navigators, 2024. URL <https://arxiv.org/abs/2406.20083>.
- [5] N. M. M. Shafiullah, C. Paxton, L. Pinto, S. Chintala, and A. Szlam. Clip-fields: Weakly supervised semantic fields for robotic memory, 2023. URL <https://arxiv.org/abs/2210.05663>.
- [6] Q. Gu, A. Kuwajerwala, S. Morin, K. M. Jatavallabhula, B. Sen, A. Agarwal, C. Rivera, W. Paul, K. Ellis, R. Chellappa, C. Gan, C. M. de Melo, J. B. Tenenbaum, A. Torralba, F. Shkurti, and L. Paull. Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning, 2023. URL <https://arxiv.org/abs/2309.16650>.
- [7] D. Maggio, Y. Chang, N. Hughes, M. Trang, D. Griffith, C. Dougherty, E. Cristofalo, L. Schmid, and L. Carlone. Clío: Real-time task-driven open-set 3d scene graphs, 2024. URL <https://arxiv.org/abs/2404.13696>.
- [8] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik. Lerf: Language embedded radiance fields. In *International Conference on Computer Vision (ICCV)*, 2023.
- [9] P. Liu, Y. Orru, J. Vakil, C. Paxton, N. M. M. Shafiullah, and L. Pinto. Ok-robot: What really matters in integrating open-knowledge models for robotics, 2024. URL <https://arxiv.org/abs/2401.12202>.
- [10] R.-Z. Qiu, Y. Hu, G. Yang, Y. Song, Y. Fu, J. Ye, J. Mu, R. Yang, N. Atanasov, S. Scherer, and X. Wang. Learning generalizable feature fields for mobile manipulation, 2024. URL <https://arxiv.org/abs/2403.07563>.
- [11] B. Bolte, A. Wang, J. Yang, M. Mukadam, M. Kalakrishnan, and C. Paxton. Usa-net: Unified semantic and affordance representations for robot memory, 2023. URL <https://arxiv.org/abs/2304.12164>.
- [12] M. Chang, T. Gervet, M. Khanna, S. Yenamandra, D. Shah, S. Y. Min, K. Shah, C. Paxton, S. Gupta, D. Batra, R. Mottaghi, J. Malik, and D. S. Chaplot. Goat: Go to any thing, 2023. URL <https://arxiv.org/abs/2311.06430>.
- [13] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard. Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation, 2024. URL <https://arxiv.org/abs/2403.17846>.
- [14] T. Gervet, S. Chintala, D. Batra, J. Malik, and D. S. Chaplot. Navigating to objects in the real world. *Science Robotics*, 8(79):eadf6991, 2023. doi:10.1126/scirobotics.adf6991. URL <https://www.science.org/doi/abs/10.1126/scirobotics.adf6991>.

- 180 [15] D. Z. Chen, A. X. Chang, and M. Nießner. Scanrefer: 3d object localization in rgb-d scans
181 using natural language, 2020. URL <https://arxiv.org/abs/1912.08830>.
- 182 [16] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-
183 annotated 3d reconstructions of indoor scenes, 2017. URL <https://arxiv.org/abs/1702.04405>.
184
- 185 [17] K. Yadav, R. Ramrakhya, S. K. Ramakrishnan, T. Gervet, J. Turner, A. Gokaslan, N. Maestre,
186 A. X. Chang, D. Batra, M. Savva, et al. Habitat-matterport 3d semantics dataset. In *Pro-*
187 *ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
188 4927–4936, 2023.
- 189 [18] G. Baruch, Z. Chen, A. Dehghan, T. Dimry, Y. Feigin, P. Fu, T. Gebauer, B. Joffe, D. Kurz,
190 A. Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding
191 using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*, 2021.
- 192 [19] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell,
193 P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from
194 natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>.
- 195 [20] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer. Sigmoid loss for language image pre-
196 training, 2023. URL <https://arxiv.org/abs/2303.15343>.
- 197 [21] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland,
198 L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár,
199 and C. Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint*
200 *arXiv:2408.00714*, 2024. URL <https://arxiv.org/abs/2408.00714>.
- 201 [22] S. Yenamandra, A. Ramachandran, M. Khanna, K. Yadav, D. S. Chaplot, G. Chhablani,
202 A. Clegg, T. Gervet, V. Jain, R. Partsey, R. Ramrakhya, A. Szot, T.-Y. Yang, A. Edsinger,
203 C. Kemp, B. Shah, Z. Kira, D. Batra, R. Mottaghi, Y. Bisk, and C. Paxton. The homerobot
204 open vocab mobile manipulation challenge. In *Thirty-seventh Conference on Neural Infor-*
205 *mation Processing Systems: Competition Track*, 2023. URL [https://aihabitat.org/](https://aihabitat.org/challenge/2023_homerobot_ovmm/)
206 [challenge/2023_homerobot_ovmm/](https://aihabitat.org/challenge/2023_homerobot_ovmm/).
- 207 [23] N. Yokoyama, A. Clegg, J. Truong, E. Undersander, T.-Y. Yang, S. Arnaud, S. Ha, D. Batra,
208 and A. Rai. ASC: Adaptive skill coordination for robotic mobile manipulation. *arXiv preprint*
209 *arXiv:2304.00410*, 2023.
- 210 [24] C. Huang, O. Mees, A. Zeng, and W. Burgard. Visual language maps for robot navigation,
211 2023. URL <https://arxiv.org/abs/2210.05714>.
- 212 [25] A. Majumdar, A. Shrivastava, S. Lee, P. Anderson, D. Parikh, and D. Batra. Improving
213 vision-and-language navigation with image-text pairs from the web. In *ECCV*, pages 259–
214 274. Springer, 2020.
- 215 [26] J. Krantz, S. Lee, J. Malik, D. Batra, and D. S. Chaplot. Instance-specific image goal navi-
216 gation: Training embodied agents to find object instances. *arXiv preprint arXiv:2211.15876*,
217 2022.
- 218 [27] M. Hahn, D. S. Chaplot, S. Tulsiani, M. Mukadam, J. M. Rehg, and A. Gupta. No rl, no sim-
219 ulation: Learning to navigate without navigating. *Advances in Neural Information Processing*
220 *Systems*, 34:26661–26673, 2021.
- 221 [28] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov. Object goal
222 navigation using goal-oriented semantic exploration. In H. Larochelle, M. Ran-
223 zato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Infor-*
224 *mation Processing Systems*, volume 33, pages 4247–4258. Curran Associates, Inc.,
225 2020. URL [https://proceedings.neurips.cc/paper_files/paper/2020/](https://proceedings.neurips.cc/paper_files/paper/2020/file/2c75cf2681788adaca63aa95ae028b22-Paper.pdf)
226 [file/2c75cf2681788adaca63aa95ae028b22-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/2c75cf2681788adaca63aa95ae028b22-Paper.pdf).

- 227 [29] N. Yokoyama, S. Ha, and D. Batra. Success weighted by completion time: A dynamics-aware
 228 evaluation criteria for embodied navigation. In *2021 IEEE/RSJ International Conference on*
 229 *Intelligent Robots and Systems (IROS)*, pages 1562–1569, 2021. doi:10.1109/IROS51168.
 230 2021.9636743.
- 231 [30] X. Zhao, H. Agrawal, D. Batra, and A. G. Schwing. The surprising effectiveness of visual
 232 odometry techniques for embodied pointgoal navigation. In *Proceedings of the IEEE/CVF*
 233 *International Conference on Computer Vision (ICCV)*, pages 16127–16136, October 2021.
- 234 [31] D. Batra, A. Gokaslan, A. Kembhavi, O. Maksymets, R. Mottaghi, M. Savva, A. Toshev, and
 235 E. Wijmans. Objectnav revisited: On evaluation of embodied agents navigating to objects.
 236 *CoRR*, abs/2006.13171, 2020. URL <https://arxiv.org/abs/2006.13171>.
- 237 [32] A. Melnik, M. Büttner, L. Harz, L. Brown, G. C. Nandi, A. PS, G. K. Yadav, R. Kala, and
 238 R. Haschke. Uniteam: Open vocabulary mobile manipulation challenge, 2023. URL <https://arxiv.org/abs/2312.08611>.
- 240 [33] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using kinect-style
 241 depth cameras for dense 3d modeling of indoor environments. *International Journal of Robotic*
 242 *Research - IJRR*, 31:647–663, 04 2012. doi:10.1177/0278364911434148.
- 243 [34] S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas. Probabilistic data association for
 244 semantic slam. In *2017 IEEE international conference on robotics and automation (ICRA)*,
 245 pages 1722–1729. IEEE, 2017.
- 246 [35] L. Zhang, L. Wei, P. Shen, W. Wei, G. Zhu, and J. Song. Semantic slam based on object
 247 detection and improved octomap. *IEEE Access*, 6:75545–75559, 2018.
- 248 [36] L. Ma, J. Stückler, C. Kerl, and D. Cremers. Multi-view deep learning for consistent seman-
 249 tic mapping with rgb-d cameras. In *2017 IEEE/RSJ International Conference on Intelligent*
 250 *Robots and Systems (IROS)*, pages 598–605. IEEE, 2017.
- 251 [37] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov. Object goal navigation using
 252 goal-oriented semantic exploration. *Advances in Neural Information Processing Systems*, 33:
 253 4247–4258, 2020.
- 254 [38] H. Ha and S. Song. Semantic abstraction: Open-world 3d scene understanding from 2d vision-
 255 language models, 2022.
- 256 [39] N. M. M. Shafiqullah, C. Paxton, L. Pinto, S. Chintala, and A. Szlam. Clip-fields: Weakly
 257 supervised semantic fields for robotic memory, 2023.
- 258 [40] C. Huang, O. Mees, A. Zeng, and W. Burgard. Visual language maps for robot navigation.
 259 In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10608–
 260 10615. IEEE, 2023.
- 261 [41] B. Chen, F. Xia, B. Ichter, K. Rao, K. Gopalakrishnan, M. S. Ryoo, A. Stone, and D. Kappler.
 262 Open-vocabulary queryable scene representations for real world planning. In *arXiv preprint*
 263 *arXiv:2209.09874*, 2022.
- 264 [42] K. M. Jatavallabhula, A. Kuwajerwala, Q. Gu, M. Omama, T. Chen, S. Li, G. Iyer, S. Saryazdi,
 265 N. Keetha, A. Tewari, et al. Conceptfusion: Open-set multimodal 3d mapping. *arXiv preprint*
 266 *arXiv:2302.07241*, 2023.
- 267 [43] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik. Lerf: Language embedded ra-
 268 diance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,
 269 pages 19729–19739, 2023.
- 270 [44] M. Ji, R.-Z. Qiu, X. Zou, and X. Wang. Graspplats: Efficient manipulation with 3d feature
 271 splatting. *arXiv preprint arXiv:2409.02084*, 2024.

- 272 [45] O. Shorinwa, J. Tucker, A. Smith, A. Swann, T. Chen, R. Firoozi, M. D. Kennedy, and
273 M. Schwager. Splat-mover: Multi-stage, open-vocabulary robotic manipulation via editable
274 gaussian splatting. In *8th Annual Conference on Robot Learning*, 2024.
- 275 [46] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time
276 radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023. URL [https://
277 //repo-sam.inria.fr/fungraph/3d-gaussian-splatting/](https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/).
- 278 [47] B. Bolte, A. Wang, J. Yang, M. Mukadam, M. Kalakrishnan, and C. Paxton. Usa-net: Unified
279 semantic and affordance representations for robot memory, 2023.
- 280 [48] Y. Wang, Z. Li, M. Zhang, K. Driggs-Campbell, J. Wu, L. Fei-Fei, and Y. Li. D3 fields:
281 Dynamic 3d descriptor fields for zero-shot generalizable robotic manipulation. *arXiv preprint
282 arXiv:2309.16118*, 2023.
- 283 [49] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: part i. *IEEE robotics
284 & automation magazine*, 13(2):99–110, 2006.
- 285 [50] Z. Song, G. Zhang, J. Xie, L. Liu, C. Jia, S. Xu, and Z. Wang. Voxelnexfusion: A simple, uni-
286 fied and effective voxel fusion framework for multi-modal 3d object detection. *arXiv preprint
287 arXiv:2401.02702*, 2024.
- 288 [51] W. Shi, J. Xu, D. Zhu, G. Zhang, X. Wang, J. Li, and X. Zhang. Rgb-d semantic segmentation
289 and label-oriented voxelgrid fusion for accurate 3d semantic mapping. *IEEE transactions on
290 circuits and systems for video technology*, 32(1):183–197, 2021.
- 291 [52] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger. Fusion++: Volumetric
292 object-level slam. In *2018 international conference on 3D vision (3DV)*, pages 32–41. IEEE,
293 2018.
- 294 [53] G. S. Krishna, K. Supriya, and S. Baidya. 3ds-slam: A 3d object detection based semantic
295 slam towards dynamic indoor environments. *arXiv preprint arXiv:2310.06385*, 2023.
- 296 [54] E. Michael, T. Summers, T. A. Wood, C. Manzie, and I. Shames. Probabilistic data association
297 for semantic slam at scale. In *2022 IEEE/RSJ International Conference on Intelligent Robots
298 and Systems (IROS)*, pages 4359–4364. IEEE, 2022.
- 299 [55] D. Maggio, M. Abate, J. Shi, C. Mario, and L. Carlone. Loc-nerf: Monte carlo localization us-
300 ing neural radiance fields. In *2023 IEEE International Conference on Robotics and Automation
301 (ICRA)*, pages 4018–4025. IEEE, 2023.
- 302 [56] A. Rosinol, J. J. Leonard, and L. Carlone. Nerf-slam: Real-time dense monocular slam with
303 neural radiance fields. In *2023 IEEE/RSJ International Conference on Intelligent Robots and
304 Systems (IROS)*, pages 3437–3444. IEEE, 2023.
- 305 [57] H. Matsuki, R. Murai, P. H. Kelly, and A. J. Davison. Gaussian splatting slam. In *Proceedings
306 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18039–
307 18048, 2024.
- 308 [58] C. Yan, D. Qu, D. Xu, B. Zhao, Z. Wang, D. Wang, and X. Li. Gs-slam: Dense visual slam
309 with 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision
310 and Pattern Recognition*, pages 19595–19604, 2024.
- 311 [59] Y. Qiu, C. Wang, W. Wang, M. Henein, and S. Scherer. Airdos: Dynamic slam benefits from
312 articulated objects. In *2022 International Conference on Robotics and Automation (ICRA)*,
313 pages 8047–8053. IEEE, 2022.
- 314 [60] L. Cui and C. Ma. Sof-slam: A semantic visual slam for dynamic environments. *IEEE access*,
315 7:166528–166539, 2019.

- 316 [61] N. Brasch, A. Bozic, J. Lallemand, and F. Tombari. Semantic monocular slam for highly
317 dynamic environments. In *2018 IEEE/RSJ International Conference on Intelligent Robots and*
318 *Systems (IROS)*, pages 393–400. IEEE, 2018.
- 319 [62] C. Yu, Z. Liu, X.-J. Liu, F. Xie, Y. Yang, Q. Wei, and Q. Fei. Ds-slam: A semantic visual slam
320 towards dynamic environments. In *2018 IEEE/RSJ international conference on intelligent*
321 *robots and systems (IROS)*, pages 1168–1174. IEEE, 2018.
- 322 [63] S. Song, H. Lim, A. J. Lee, and H. Myung. Dynavins: a visual-inertial slam for dynamic
323 environments. *IEEE Robotics and Automation Letters*, 7(4):11523–11530, 2022.
- 324 [64] P. Yu, C. Guo, y. Liu, and H. Zhang. Fusing semantic segmentation and object detection for
325 visual slam in dynamic scenes. In *Proceedings of the 27th ACM Symposium on Virtual Reality*
326 *Software and Technology*, pages 1–7, 2021.
- 327 [65] B. Bescos, J. M. Fácil, J. Civera, and J. Neira. Dynaslam: Tracking, mapping, and inpainting
328 in dynamic scenes. *IEEE Robotics and Automation Letters*, 3(4):4076–4083, 2018.
- 329 [66] L. Schmid, M. Abate, Y. Chang, and L. Carlone. Khronos: A unified approach for spatio-
330 temporal metric-semantic slam in dynamic environments. In *Proc. of Robotics: Science and*
331 *Systems*, 2024.
- 332 [67] J. C. Virgolino Soares, V. S. Medeiros, G. F. Abati, M. Becker, G. Caurin, M. Gattass, and
333 M. A. Meggiolaro. Visual localization and mapping in dynamic and changing environments.
334 *Journal of Intelligent & Robotic Systems*, 109(4):95, 2023.
- 335 [68] B. Bescos, C. Campos, J. D. Tardós, and J. Neira. Dynaslam ii: Tightly-coupled multi-object
336 tracking and slam. *IEEE robotics and automation letters*, 6(3):5191–5198, 2021.
- 337 [69] M. Henein, J. Zhang, R. Mahony, and V. Ila. Dynamic slam: The need for speed. In *2020*
338 *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2123–2129. IEEE,
339 2020.
- 340 [70] D. F. Henning, T. Laidlow, and S. Leutenegger. Bodyslam: Joint camera localisation, mapping,
341 and human motion tracking. In *European Conference on Computer Vision*, pages 656–673.
342 Springer, 2022.
- 343 [71] J. Haviland, N. Sünderhauf, and P. Corke. A holistic approach to reactive mobile manipulation.
344 *IEEE Robotics and Automation Letters*, 7(2):3122–3129, 2022.
- 345 [72] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakr-
346 ishnan, K. Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances.
347 *arXiv preprint arXiv:2204.01691*, 2022.
- 348 [73] Y. Du, D. Ho, A. Alemi, E. Jang, and M. Khansari. Bayesian imitation learning for end-to-end
349 mobile manipulation. In *International Conference on Machine Learning*, pages 5531–5546.
350 PMLR, 2022.
- 351 [74] J. Wong, A. Tung, A. Kurenkov, A. Mandlekar, L. Fei-Fei, S. Savarese, and R. Martín-Martín.
352 Error-aware imitation learning from teleoperation data for mobile manipulation. In *Conference*
353 *on Robot Learning*, pages 1367–1378. PMLR, 2022.
- 354 [75] S. Uppal, A. Agarwal, H. Xiong, K. Shaw, and D. Pathak. Spin: Simultaneous perception,
355 interaction and navigation, 2024. URL <https://arxiv.org/abs/2405.07991>.
- 356 [76] O. Team. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- 357 [77] G. T. Google. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of
358 context, 2024. URL <https://arxiv.org/abs/2403.05530>.

- 359 [78] J. Yang, X. Chen, S. Qian, N. Madaan, M. Iyengar, D. F. Fouhey, and J. Chai. Llm-grounder:
360 Open-vocabulary 3d visual grounding with large language model as an agent, 2023. URL
361 <https://arxiv.org/abs/2309.12311>.
- 362 [79] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. Vqa: Visual
363 question answering. In *Proceedings of the IEEE international conference on computer vision*,
364 pages 2425–2433, 2015.
- 365 [80] A. Majumdar, A. Ajay, X. Zhang, P. Putta, S. Yenamandra, M. Henaff, S. Silwal, P. Mcvay,
366 O. Maksymets, S. Arnaud, et al. Openeqa: Embodied question answering in the era of founda-
367 tion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
368 *Recognition*, pages 16488–16498, 2024.
- 369 [81] M. Minderer, A. Gritsenko, and N. Houlsby. Scaling open-vocabulary object detection, 2024.
370 URL <https://arxiv.org/abs/2306.09683>.
- 371 [82] N. Yokoyama, S. Ha, D. Batra, J. Wang, and B. Bucher. Vlfm: Vision-language frontier maps
372 for zero-shot semantic navigation. In *2024 IEEE International Conference on Robotics and*
373 *Automation (ICRA)*, pages 42–48. IEEE, 2024.
- 374 [83] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu. Anygrasp:
375 Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on*
376 *Robotics*, 2023.




Figure 4: Statistics of failure, broken down by failure modes, in our real robot experiments. Statistics are collected over three environments and 30 open-vocabulary pick-and-drop queries on objects whose locations change over time.

377 5 Related Works

378 5.1 Open Vocabulary Mobile Manipulation (OVMM)

379 Navigating to arbitrary goals in open ended environments and manipulating them has become a key
 380 challenge in robotic manipulation [22, 23]. This line of query follows Open-Vocabulary Navigation
 381 systems [5, 24], which builds upon prior object and point goal navigation literature [14, 25, 26, 27,
 382 28, 29, 30, 31, 12] which attempted navigation to points, or fixed set of objects and object categories.
 383 OVMM is a naturally harder challenge as it requires an ability to handle arbitrary queries, and
 384 “navigation to manipulation” transfer – which means unlike pure navigation, the robot needs to get
 385 close to the environment objective and obstacles. In the OVMM challenge [22], modular solutions
 386 such as [1, 32, 13] outperformed the competition. More recently, OK-Robot [9] performed extensive
 387 real-world evaluations of the challenges in OVMM and demonstrated a system that achieves 58.5%
 388 success rate in static home environments. We extend this work by enabling manipulation in changing
 389 environments.

390 5.2 Spatio-semantic Memory

391 Early works in spatio-semantic memory [33, 34, 35, 36, 37] created semantic maps for limited
 392 categories based on mostly ad-hoc deep neural networks. Later work builds upon representations
 393 derived from pre-trained vision language models, such as [38, 39, 40, 41, 42, 43, 6, 7]. These works
 394 use a voxel map or neural feature field as their map representation. Some recent models [44, 45]
 395 have used Gaussian splats [46] to represent semantic memory for manipulation. Most of these
 396 models show object localization in pre-mapped scenes, while CLIP-Fields [5], VLMs [24], and
 397 NLMap-SayCan [41] show integration with real robots for indoor navigation tasks. Some recent
 398 works [47, 48, 10] extend this task to include an affordance model or manipulation primitives. Our
 399 work builds upon the voxel map based spatio-semantic memory literature and extends them to dy-
 400 namic environments where both objects and obstacles can change over time.

401 5.3 Mapping and Navigating Dynamic Environments


402 For robot navigation, Simultaneous Localization and Mapping (SLAM) [49] methods are crucial.
 403 However, practical SLAM instances based on voxels [50, 51], objects [52, 53], landmark [34, 54],
 404 NeRF [55, 56], and Gaussian splats [57, 58] tend to make the simplifying assumption that the world



Kitchen



Game room



Meeting room

Figure 5: Real robot experiments in three different environments: kitchen, game room, and meeting room. In each environment, we modify the environment thrice and run 10 pick-and-drop queries.




Figure 6: DynaMem keeps its memory stored in a sparse voxel grid with associated information at each voxel.




Figure 7: Updating DynaMem by adding new points to it, alongside the update rules for the stored information.

405 is static. Some sparse SLAM methods improve on dynamic environments by estimating underlying
 406 state [59, 60, 61, 62, 63, 64, 65, 66, 67] or explicitly modeling moving objects [68, 69, 70]. Some
 407 methods also forego a map and rely on reactive policies to navigate dynamic environments [71, 72,

408 [73, 74, 75], although they generally tackle local movement and not global navigation. Our work
409 relies on SLAM systems that are stable under environment dynamics, and focuses on building a
410 dynamic semantic memory based off of online exploration and observations.

411 6 Querying DynaMem for Object Localization

412 As described in Section 2.1, we define the object localization or 3D visual grounding problem as
413 a function mapping a text query and posed RGBD images to either the 3D coordinate of the query
414 object, or \emptyset if the object is not in the scene. Unlike previous work, we abstain from returning a
415 location when an object is not found. To enable this, we factor this grounding problem into two sub-
416 problems. The first is finding the latest image where the queried object could have appeared. The
417 second is identifying whether the object is actually present in that image. For the first sub-problem,
418 we propose two alternate approaches of visual grounding: one using the intrinsic semantic features
419 of DynaMem, and another using state-of-the-art multimodal LLMs such as GPT-4o [76] and Gemini
420 1.5 Pro [77].

421 **Embedded Vision Language Features:** Vision Language Models (VLMs) such as CLIP [19] and
422 SigLIP [20] possess an ability to embed both images and languages into the same latent space, where
423 the similarity between an image and a text object can be calculated by simply taking the dot product
424 between the two latent representation vectors. We use this property of the embedding vectors to
425 query our voxel map with open-vocabulary text queries.

426 As described in Section 2.2, we convert the incoming images to point-wise image features, and
427 embed them into our voxels. When we have a new language query, we calculate its latent embedding
428 using the VLM text encoder, and find the voxel whose feature has the highest dot product with the
429 text embedding. Once we find the right voxel, we simply retrieve its associated latest image from
430 our data structure as shown in Fig. 3.

431 As a bonus feature, we can also return $n > 1$ possible objects for a single query. We do this by using
432 a DBSCAN clustering of voxels similar to [78], and returning the images associated with the most
433 aligned voxel in top- n clusters.

434 **Multimodal Large Language Models (mLLMs):** We note that the problem of finding the latest
435 image where an object may appear is similar to the problem of visual question-answer (VQA) [79].
436 Since we fully rely on pretrained models to build our map, we pose this multi-image VQA problem
437 as an mLLM QA problem similar to OpenEQA [80].

438 We show in Fig. 8 how we query the mLLMs to solve the visual grounding query. We give the
439 model a sequence of our latest environment observations images and ask the model for the index of
440 the last image where the queried object was observed. We additionally instruct the model to respond
441 “None” if the object was not observed in any image. Note that, unlike OpenEQA [80], we only pass
442 the RGB images to the mLLM, and not the depth or camera pose. Similarly, we only ask for an
443 image index, and not a full textual answer.

444 **Handling Absence of Object:** Several previous methods [5, 8, 9] assume that the queried object is
445 always present in the scene, and always responds with the object that is the best match to the query.
446 However, this often results in high false-positive failure cases. For example, in a scene with no red
447 cups and a blue cup, the method may respond with the location of the blue cup in response to the
448 query “red cup”.

449 For this reason, we locate objects in two stages. First, we find the best candidate image where the
450 object may have been seen (Section 6). Then, we use an open-vocabulary object detector model such
451 as OWL-v2 [81] to search that image for the queried object (Fig. 3). If we don’t find the queried
452 object, we assume that the object has either moved, or the response from the voxelmap or mLLM
453 was inaccurate, and respond with “object not found”. If OWL-v2 returns an object bounding box,
454 we find the median pixel from the object mask and return its 3D location.

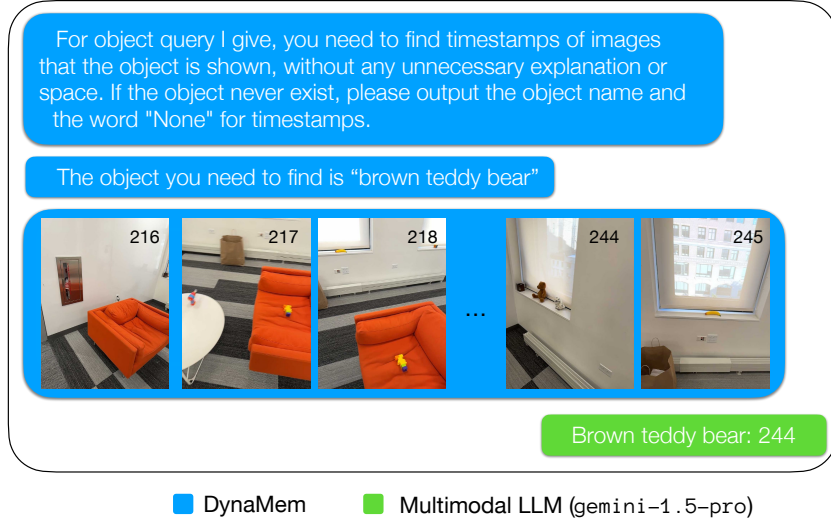


Figure 8: The prompting system for querying multimodal LLMs such as GPT-4o or Gemini-1.5 for the image index for an object query.

455 One important hyperparameter for this mLLM query is the maximum number of images included in
 456 the prompt. Longer context needs longer processing time and potentially includes outdated informa-
 457 tion, while short context might not include all information and thus will miss objects. We optimize
 458 the context by excluding completely outdated images: all images I with no voxel pointing to them
 459 are deleted. This filtering increases mLLM context utilization. We set Gemini as our base model
 460 and 60 as our query image limit since Gemini context can fit 60 images, which is twice as many as
 461 GPT-4o.

462 7 Robot Navigation and Exploration

463 To navigate in a real-world environment, robots use an obstacle map in conjunction with a navigation
 464 algorithm like A* in [24, 9]. We use a simple voxel-projection strategy to build an obstacle map. Due
 465 to the depth observation noise, we simply set a threshold for the ground (0.2m for our experiments),
 466 and project all the voxels above that z -threshold as the obstacles in our map. The voxels below the
 467 threshold are projected into the 2D obstacle map as navigable points. Finally, the points in the map
 468 that are not marked as either obstacle or navigable are marked as explorable points.

469 **Exploration Primitives:** Since our robot does not start with an environment map, it explores the
 470 environment with frontier based methods to build the map. We can further accelerate this process
 471 by providing exploration guidance. Based on the current status of the map, DynaMem provides an
 472 exploration value function to accelerate the exploration process both for building and updating the
 473 map.

474 We provide two value-based exploration maps: one time-based, and one semantic-similarity-
 475 based [82]. The time-based value map prioritizes the least-recently seen points. If the current time
 476 is T , and the last-seen time of voxel (x, y, z) is $t_{x,y,z}$, the temporal value map \mathbb{V}_T is expressed as:

$$\mathbf{T}^*[x, y] = \max_z (T - t_{x,y,z})$$

$$\mathbb{V}_T[x, y] = \sigma(-\beta_T(\mathbf{T}^*[x, y] - \mu_T))$$

477 where β_T, μ_T are hyper-parameters and σ is the sigmoid function. Similarly, if the VLM feature at
 478 voxel (x, y, z) is $f_{x,y,z}$, and the VLM feature for the language query is f_q , then the similarity-based

479 value map \mathbb{V}_S is be expressed as:

$$\mathbf{S}^*[x, y] = \max_z(f_q \cdot f_{x,y,z})$$
$$\mathbb{V}_S[x, y] = \sigma(-\beta_S(\mathbf{S}^*[x, y] - \mu_S))$$

480 where once again β_S, μ_S are hyperparameters. We may also linearly combine $\mathbb{V}_T, \mathbb{V}_S$ to balance
481 our exploration between last seen time and semantic similarity.

482 Finally, since the environment may be dynamic, we convert our navigation algorithm from open-loop
483 to closed-loop. The robot, instead of executing the entire navigation plan generated by A*, stops
484 after the first seven waypoints (approx. 0.7 to 1 meters). Then, the robot scans the environment,
485 updates the map, and moves according to a new plan. The robot repeats these steps until its distance
486 to the target is lower than a predefined threshold.

487 8 Ablations on an Offline Benchmark

488 Running real robot OVMM experiments can be expensive and time-consuming. So, we developed
489 an offline benchmark called DynaBench to easily evaluate dynamic 3D visual grounding algorithms
490 on dynamic environments and perform algorithmic ablations. The benchmark isolates the query-
491 response part of the dynamic semantic memory without robot navigation, exploration, and manipu-
492 lation.

493 8.0.1 Data Collection

494 In the real world, the robot collects its own map-building data by exploring the environments. Fol-
495 lowing this, we collect the robot’s runtime sensor data from three environments. To further enrich
496 our benchmark, we simulate this process by taking posed RGB-D images on an iPhone Pro in six
497 more environments. In all cases, we emulate environment dynamics by moving objects and obstacle
498 locations in three successive rounds.

499 8.0.2 Data Labelling and Evaluation

500 We manually annotate queries and responses in the dataset. Each query has an associated natural
501 language label q , object location $\vec{X} = (x, y, z)$, and an object radius ϵ . Since the environment is
502 dynamic, each query also has an associated time t . For evaluation, at time t (i.e. after the memory
503 algorithm has observed all the input data points with timestamp $< t$), we query the model with q . If it
504 predicts an object location $\vec{X}' = (x', y', z')$, it’s a success if $\|\vec{X} - \vec{X}'\|_2 \leq \epsilon$ and a failure otherwise.
505 Since the robot may also encounter queries for objects it has not observed yet, we emulate negative
506 queries by adding queries for objects (a) that have not been observed yet, or (b) that have been
507 observed but were subsequently removed. For both of these query types, the model must respond
508 with *not found*; otherwise it’s counted as a failure.

509 8.0.3 Evaluation Results

510 Using our offline benchmark, we ablate design decisions of DynaMem as discussed in Section 2.
511 Among these design decisions, the primary are: using feature embedding-based vs. mLLM-QA
512 based language grounding, ablating components such as point removal or abstention from the
513 algorithm, and trying different mLLMs. Due to API costs, we only evaluate Gemini models on the
514 benchmark. We present our results in Table 1.

515 We see that performance of VLM-features and mLLM-QA follows the same order in the real world
516 in the benchmark, corroborating the benchmark design. The best design choices are to both add and
517 remove points, and to cross check with OWL-v2 on top of similarity thresholding for VLM-feature
518 based grounding. For mLLM-QA based grounding, Gemini Pro outperforms Gemini Flash, and
519 voxelmap based image filtering benefits the method.

Table 1: Ablating the design choices for our query methods for DynaMem on the offline DynaBench benchmark. We also present results from five human participants to ground the performances.

Query type	Variant	Success rate
Human	(average over five participants)	81.9%
VLM-feature	default (adding and removing points)	70.6%
	only adding points	67.8%
	no OWL-v2 cross-check	59.2%
	no similarity thresholding	66.8%
mLLM-QA	default (Gemini Pro 1.5)	67.3%
	Gemini Pro 1.5, no voxelmap filtering	66.8%
	Gemini Flash 1.5	63.5%

520 9 Experiment Setup

521 We evaluate DynaMem and its impact on open-vocabulary mobile manipulation in three real-world
522 dynamic environments (Fig. 5). In each environment, we set up multiple objects as potential manip-
523 ulation targets, change the environment in three rounds, and execute 10 pick-and-drop queries over
524 the rounds We use the Hello Stretch SE3 as our mobile robot platform, and use its head-mounted
525 Intel RealSense D435 RGB-D camera to collect the input data.

526 To build a complete pick-and-drop system around DynaMem, we follow the system architecture in
527 OK-Robot [9]. In particular, we use the AnyGrasp [83] based open-vocabulary grasp system and
528 use the heuristic based dropping system. However, we use DynaMem’s exploration primitives let
529 the robot build the map of the environment and allow the robot to explore when an object is not
530 found in the memory.