

APPENDIX A
SUPPLEMENTARY DETAILS: PERPETUA*

A.1 Perpetua* Algorithm

In Algorithm 1 we present the update and prediction routines for a single receptacle in Perpetua*, omitting the receptacle index k for clarity. During the update step, we compute the likelihood, conditional evidence, and posterior weights for both the persistence and emergence models. In the prediction step, we generate the Perpetua* prediction at an arbitrary query time using our proposed Bayesian model selection criterion.

Algorithm 1: Predict and Update Functions for Perpetua* (Single Receptacle)

Parameters: Models $m \in \{\mathcal{M}_P, \mathcal{M}_E\}$, Decay rate α_0 , Forgetting factor γ for receptacle k

▷ Updates the state of each model

Update Input: Observation tuple $(y_{t_{N+1}}, t_{N+1})$

```

1 Function Update( $y_{t_{N+1}}, t_{N+1}$ ):
2   for  $m \in \{\mathcal{M}_P, \mathcal{M}_E\}$  do
3     Apply forgetting factor  $\gamma$  and update likelihood
        $p(\mathcal{Y}_{1:N+1} | x_{N+1}, m)$  via (3);
4     for  $l \leftarrow 1$  to  $L$  do
5       Update cond. evidence  $p(\mathcal{Y}_{1:N+1} | c_l, m)$ ;
6       Update posterior weights  $w_l$  via (5);
7    $N \leftarrow (N + 1)$ 

```

▷ Predicts persistence state at time t

Predict Input: Query time t for receptacle k

```

8 Function Predict( $t$ ):
9   for  $m \in \{\mathcal{M}_P, \mathcal{M}_E\}$  do
10    Evaluate predictions  $p(x_t | c_{l^*}, \mathcal{Y}_{1:N}, m)$  via (6);
11     $\alpha(t) \leftarrow \exp(-\alpha_0(t - t_N))$ ;
12    Compute model posterior  $p(\mathcal{M} | \mathcal{Y}_{1:N})$  via (8);
13    Compute final prediction  $p(x_t | \mathcal{Y}_{1:N})$  via (9);
14   return  $p(x_t | \mathcal{Y}_{1:N})$ ;

```

Tool 1: Semantic Object Search

Input: query_text (string)

Output: List of $\langle \text{object_id: string, score: float} \rangle$ or None

Description: Performs an open-vocabulary search over PredictiveGraphs using CLIP. It compares the embedding of the query_text against the visual embeddings of all mapped objects, returning the object_ids with the highest cosine similarities.

Tool 2: Predict Object Receptacle

Input: object_id (string), query_time (float)

Output: Most likely receptacle name (string), probability distribution over receptacles (list of floats)

Description: Queries the temporal object map (PredictiveGraphs) using Perpetua* to predict the location of the target object at the specified query_time. The tool applies a confidence threshold (default $\delta = 0.2$) to filter receptacles. If no receptacle exceeds the threshold, the object is reported as absent.

Tool 3: Navigate to Receptacle

Input: receptacle_id (string), object_id (string)

Output: Success status (boolean), final observation

Description: Navigates to the receptacle_id while actively searching for the target object_id. During navigation, the agent continuously: (1) executes path planning toward the receptacle, (2) updates PredictiveGraphs with RGBD observations, and (3) checks whether the target object is in view. Navigation terminates when either the object is found or the target receptacle is reached.

APPENDIX B
SUPPLEMENTARY & IMPLEMENTATION DETAILS

B.1 LLM Tools

In this section, we provide the formal specifications for the three core tools that enable embodied planning in PredictiveGraphs: *semantic search* (Tool 1), *location prediction* (Tool 2), and *active navigation* (Tool 3). For each tool, we define the required inputs, the expected outputs, and the underlying logic used to interface with PredictiveGraphs.

B.2 LLM Prompt - Switching Prior

We present the full LLM prompt used to query the LLM-based switching prior described in Sec VI-A1. We employ a few-shot prompting strategy to leverage the semantic knowledge of gemini-2.5-pro, enabling it to infer probable weekly schedules for objects based on their household context and potential usage patterns. The prompt is shown in Prompt 1

Prompt 1: LLM Prior for Perpetua*

Consider a household environment composed of two parents and two children. The environment contains various objects located in different places, where objects can appear when needed for an activity (e.g., plates for eating) or disappear when taken by a member (e.g., keys/bags when leaving). During weekdays, the parents typically leave for work in the morning and return in the evening (07:30 - 18:00), while the children attend school during the day (08:00 - 17:00). On weekends, the family tends to stay home more often, engaging in activities such as cooking, cleaning, and leisure time together.

Task: Generate a weekly schedule (Mon-Sun) for the specific object at the specific location.

Output Rules:

- Output ONLY a Python list of strings.
- Format: ['Day: HH:MM - Day: HH:MM', 'Day: HH:MM - Day: HH:MM', ...]
- Use 3-letter day abbreviations (Mon, Tue, Wed, Thu, Fri, Sat, Sun).

Example 1:

Input:

Object: Plates

Location: On Dinner Table

Type: Active Location (appear when used)

Output:

Schedule: ['Mon: 19:00 - Mon: 19:30', 'Tue: 19:00 - Tue: 19:30', ..., 'Sun: 18:00 - Sun: 18:30']

Current task:

Input:

Object: {object_name}

Location: {object_location}

Type: {object_type}

Output:

Schedule: [...]




Fig. 8. **Weekly Test Schedule:** One-week snapshot of the test set with out-of-distribution dynamics for the different objects used in the Perpetua* evaluation presented in Sec VI-A. Note how Friday and Monday follow the same dynamics as the weekend, emulating a “long weekend”.

- 1) **FreMEEn Prior:** Based on (7), this variant computes 1000 Fourier coefficients and selects the optimal subset via a held-out validation set of one week.
- 2) **LLM Prior:** Generated via gemini-2.5-pro [15] using the formulation in (7). The full prompt is shown in Appendix B.2
- 3) **Oracle Prior:** Uses step functions that perfectly match the ground truth dynamics of the training set.

B.3 Perpetua* Implementation Details

In this section, we detail the parameters and modeling choices for Perpetua*.

a) **Hyperparameters:** The prior p_t in (1) is modeled as a mixture of log-normal distributions. We set the forgetting factor $\gamma = 0.99$, which maintains an effective memory of approximately 250 observations. The decay rate is set to $\alpha_0 \approx 0.01$, causing the model to revert almost entirely to the prior after approximately 300 time units.

b) **Model Selection:** Following the configuration in Perpetua [12], we search for the optimal number of mixture components (up to a maximum of 5) for both persistence and emergence filters using the Akaike information criteria (AIC) [41]. The observation noise parameters (P_M, P_F) are estimated directly from dataset statistics.

c) **Switching Priors:** We evaluate Perpetua* with three variants of the switching prior $f(t)$:

B.4 Schedule

While the training set used to train Perpetua* in Sec VI-A reflects a standard weekly routine with distinct weekday and weekend patterns, the test set introduces a distributional shift to evaluate adaptation. Specifically, we construct a “long weekend” scenario where Monday and Friday exhibit weekend-like dynamics. This modification forces the model to adjust to sudden schedule changes, serving as a stress test for real-time adaptation. A snapshot of these shifted dynamics is provided in Fig. 8.

APPENDIX C

SUPPLEMENTARY RESULTS: PERPETUA*

C.1 Complexity Analysis and Illustrative Example: Perpetua*

In Fig. 9, we report the results of a complexity analysis of Perpetua* versus Perpetua. Perpetua* is more computationally




Fig. 9. **Computational Complexity:** prediction (top) and update (bottom) steps for *Perpetua** with a FreMEN prior versus *Perpetua*. Top: prediction time as a function of Fourier coefficients in *Perpetua** and simulation steps in *Perpetua*. Bottom: update time with up to 1 million samples. Results are averaged over five random seeds. *Perpetua** exhibits faster computation in both cases by replacing *Perpetua*’s state machine with our Bayesian model selection mechanism.

TABLE II
MEMORY FOOTPRINT OF PERPETUA*

	Quantity	Value
(a) Per mixture component	Base cost	27.89 KB
	Cost per component	0.047 KB
(b) Per edge	Base cost	0.000 KB
	Cost per edge	28.13 KB

(a) Per-edge cost as a function of mixture components. (b) Total cost as a function of edges using mixture models with five mixture components and a FreMEN prior with 1,000 coefficients.

efficient at both prediction and belief updates due to its switching mechanism based on Bayesian model selection, which avoids *Perpetua*’s reliance on expensive simulation steps.

Table II shows the memory footprint of adding extra mixture components or edges to the graph. The per-mixture cost is negligible (0.047 KB), while the per-edge cost is 28 KB for a model with five mixture components and a FreMEN prior with a thousand coefficients. New edges are only added when a new object-receptacle pair is observed, so memory scales linearly with the number of discovered relationships rather than with deployment time. For example, in a scene with 250 objects appearing in up to 4 different locations each (1,000 edges), the total overhead is **28 MB** on top of the base scene graph, which can itself reach several GB due to stored object point clouds. Overall, the additional memory and compute overhead of *Perpetua** is negligible compared to modern mapping approaches. Furthermore, in Fig. 10 we present an illustrative example showing how the different components (mixture models, Bayesian model selection) interact to yield the final *Perpetua** prediction even in the absence of data.

APPENDIX D

SIMULATION DETAILS: PREDICTIVEGRAPHS

D.1 Hierarchical Scheduling Model

The spatial distribution of each pickupable object over time is governed by a hierarchical temporal model inspired by




Fig. 10. **Illustrative Example *Perpetua**:** An object undergoes semi-static changes at varying rates (top row). We maintain a mixture of persistence and emergence filters (each comprising three components in this example) that *adapt* their belief as new observations arrive. The second row shows the dominant component of the persistence and emergence filters (see (6)) at each time step. As time progresses, our Bayesian model selection criterion (see (8)) identifies the mixture that best explains the observed data (third row). Notably, between times t_a and t_b , where no observations are available, the log posterior odds keep evolving due to our switching prior, whereas the dominant mixture component remains frozen. The bottom row shows the final *Perpetua** prediction, obtained by mixing the output of both mixture models at each time step, as shown in (9). By leveraging the switching prior (cf. Fig. 2 and Fig. 4), *Perpetua** maintains accurate predictions even in the absence of observations.

human activity patterns. For each pickupable, we generate a schedule that specifies which receptacle contains that object at each timestep. This schedule uses a hierarchical structure with different temporal scales (e.g., year $\rightarrow$ month $\rightarrow$ week $\rightarrow$ day $\rightarrow$ hour) and different divisions for each hierarchical level. Transitions between hierarchical levels and receptacles follow a Markov process. For instance, a schedule whose hierarchical levels are day and hour, with 5 different day-level patterns, will have a one day-level transition matrix and 5 different hour-level transition matrices. Following Kurenkov et al. [23], the transition matrices are defined using weights computed from the relationship between “pickupable” and “receptacle” objects defined in ProcTHOR. Duration in each receptacle is sampled from a uniform distribution, also informed by the ProcTHOR semantic prior. This hierarchical structure allows us to capture realistic patterns: objects stay in certain locations for hours or days, but may transition to different receptacles across longer timescales³. Figure 11 shows a two-week generated schedule. The first week shows standard train-time dynamics, while the second week exhibits shifted test-time dynamics, as used during our adaptation experiments in Sec. VI-B4.

³The source code will be made publicly available soon.




Fig. 11. **Dynamic Shift in ProcTHOR Adaptation Experiments:** Illustration of the train and test time dynamics for the adaptation experiments in Sec VI-B4. The first week displays standard training dynamics for the “newspaper” object across three receptacles. In the second week, these dynamics change, introducing a shift that challenges predictive methods, requiring approaches capable of real-time adaptation such as Perpetua*.

D.2 Temporal Resolution

Observations are collected at a fixed temporal δt of 1 hour. We generate continuous trajectory data by simulating 4 weeks of time per environment, yielding a total of 672 timesteps per environment (partitioned into two weeks for training, one for validation, and one for testing). We use 15 different ProcThor scenes, which we select to have at least 5 rooms each and a navigable space between 100 m² and 150 m². A top-down view of the scenes is shown in Fig. 16

D.3 Relationship Between Groundtruth and Privileged Data

The groundtruth assignment indicates the simulator’s “true” state (the receptacle that contains the object), while the privileged information filters this through the current observability constraints imposed by the camera viewpoint and rendering. Formally:

$$\text{privileged}_{s,r,t} = \begin{cases} 1 & \text{if } \text{gt}_{s,r,t} = 1 \wedge \text{obj. visible,} \\ 0 & \text{if } \text{gt}_{s,r,t} = 0 \wedge \text{rec. visible,} \\ -1 & \text{if rec. hidden,} \\ -2 & \text{if } r \text{ invalid for } p; \end{cases} \quad (11)$$

where 1 denotes presence, 0 absence, -1 missing data, and -2 that the receptacle r is invalid for object s .

APPENDIX E

REAL-WORLD DETAILS: PREDICTIVEGRAPHS

E.1 Real-World Details

For our real-world evaluation, we construct a dataset where we manually relocate objects according to a known ground-truth schedule. We deployed an Agilex Ranger Mini 2.0 robot equipped with an Intel RealSense D435 camera and teleoperated it through multiple trajectories to gather RGB-D observations. RTAB-Map [40] was used as the pose estimation backend and mapping tool. To




Fig. 12. **Real-World Environment:** (Top) Top-down view showing receptacle and door bounding boxes. Receptacles are color-coded by room: pink (first room), teal (second room), and mustard (third room). Door bounding boxes are highlighted in blue. (Bottom) An isometric view of the three-room laboratory setup used for all real-world experiments.

evaluate PredictiveGraphs under realistic perception noise, we generated object-receptacle training labels for Perpetua* using a custom perception pipeline based on SAM3 [42] and the association function described in Sec V-A

The environment consists of a three-room laboratory connected by two doors (Fig. 12), containing a total of 28 receptacles distributed as 10, 10, and 8 across the rooms, respectively. Throughout the experiments, the two doors were periodically opened and closed to vary the environment’s topological connectivity. Scans were taken bi-hourly from 9:00 AM to 5:00 PM over a span of three weeks, yielding a dataset of 84 scans. For each scan, we provide RGB-D frames and camera poses in a common origin frame, along with manually annotated 3D bounding boxes for the receptacles.

During data collection, we used the following 21 objects: greenbowl, bell, panda, penguin, redbowl, bluecup, yellowcup, greencube, bluecube, headphones, duck, egg, carrot, hairbrush, redbottle, redmug, redplate, woodenspoon, spoon, spatula, blueball.

E.2 Hardware

All simulation experiments were conducted on an institutional cluster equipped with Nvidia RTX 8000 and L40S

TABLE III
ADAPTIVE NAVIGATION EXPERIMENT: DETAILED VERSION OF FIG 6 (LEFT).

Perception	Method	Static Setting (80)		Dynamic Setting (45)				Overall (125)	
		Presence (80)		Presence (38)		Absence (7)		Average	
		Success (%) $\uparrow$	SPL (%) $\uparrow$	Success (%) $\uparrow$	SPL (%) $\uparrow$	AR (%) $\uparrow$	Dist (m) $\downarrow$	Success (%) $\uparrow$	SPL (%) $\uparrow$
Privileged	CG [16]	52.0 $\pm$ 9.9	51.6 $\pm$ 9.8	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	66.7 $\pm$ 19.2	0.97 $\pm$ 0.53	37.6 $\pm$ 8.1	36.3 $\pm$ 8.4
	Ours (FreMEEn [13])	79.3 $\pm$ 2.6	78.1 $\pm$ 3.0	53.0 $\pm$ 13.4	48.8 $\pm$ 11.9	44.4 $\pm$ 29.4	9.3 $\pm$ 2.7	72.0 $\pm$ 3.6	70.2 $\pm$ 2.8
	Ours	89.3 $\pm$ 2.7	83.4 $\pm$ 1.3	65.4 $\pm$ 10.4	56.3 $\pm$ 9.9	88.9 $\pm$ 11.1	7.7 $\pm$ 2.3	83.2 $\pm$ 4.3	75.9 $\pm$ 2.9
CLIP + LLM	CG [16]	47.4 $\pm$ 8.5	47.0 $\pm$ 8.4	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	66.7 $\pm$ 19.2	0.97 $\pm$ 0.53	34.4 $\pm$ 7.4	33.1 $\pm$ 7.7
	Ours (FreMEEn [13])	57.1 $\pm$ 4.8	55.9 $\pm$ 4.9	37.2 $\pm$ 9.4	36.2 $\pm$ 8.7	44.4 $\pm$ 29.4	10.9 $\pm$ 3.8	52.0 $\pm$ 5.7	50.7 $\pm$ 5.2
	Ours	58.5 $\pm$ 3.6	57.3 $\pm$ 3.8	38.6 $\pm$ 8.7	37.2 $\pm$ 8.1	100.0 $\pm$ 5.1	9.8 $\pm$ 4.3	55.4 $\pm$ 5.1	51.6 $\pm$ 4.4

The adaptive capabilities of Perpetua* enable better performance even under distributional shifts in environment dynamics. While perception errors degrade performance by introducing spurious belief updates, Perpetua* maintains its overall advantage due to its adaptation capabilities.




Fig. 13. **Qualitative Navigation Results in ProcTHOR:** The top row shows a navigation sequence where PredictiveGraphs is queried 46 hours after the last map update (top-down map shown in the first row, third column of Fig. 16). In the first column, our method predicts that the soap bottle has moved from the dresser (red spline tip) to a side table (green spline tip). The second column shows the robot navigating to the predicted location, with its path marked in blue. The final column displays the successful detection of the soap bottle. The bottom row depicts a similar example with an alarm clock queried 72 hours after the last map update (top-down map shown in the third row, first column of Fig. 16).

GPUs. The Agilex Ranger Mini 2.0 was equipped with an Nvidia Jetson Orin for onboard control, but model inference was performed offboard, on an Nvidia RTX 2080. For both the LLM-based navigation agent and the LLM-based predictive baselines we use gpt-5.2.

APPENDIX F EXTENDED REAL-WORLD RESULTS

F.1 Simulation results

Table III presents the comprehensive adaptive navigation results, expanding on the results provided in Sec. VI-B4. Here, we ablate PredictiveGraphs by comparing the fully-predictive FreMEEn estimator against our adaptive Perpetua* method. Under privileged perception, PredictiveGraphs with Perpetua* achieves superior navigation performance, effectively adapting to the shift in test-time dynamics. Even when

perception errors degrade performance, particularly in the dynamic setting, Perpetua* maintains its advantage over the FreMEEn variant. As demonstrated by the privileged results, we expect that as perception pipelines improve, the performance of PredictiveGraphs will scale accordingly.

Additionally, Fig. 13 presents qualitative results for the ProcTHOR navigation experiments from Sec. VI-B4. This figure shows how PredictiveGraphs accurately predicts the location of target objects at different query times. This capability allows the agent to exploit predictive information to successfully navigate to its target, even when the underlying scene graph has not been updated for several hours.

F.2 Real-World results

In this section, we provide additional qualitative results to supplement the real-world navigation experiments discussed in Sec. VI-B5. Figure 14 presents an extended version of




Fig. 14. **Qualitative Navigation Result in Real-World - Topological Graph Change:** Extended version of Fig. 7. The open-vocabulary query is “Is there something useful to comb my hair?”, with query time 146 hours after the last map update. (1) Shows the last known position of the relevant object identified by PredictiveGraphs (a hair brush). (2) Our method predicts that the hair brush is in another room and detects that the door along the shortest path is closed; this updates the cost map, prompting the planner to generate a longer but feasible path. (3) The agent successfully locates the target object.

Fig. 7 demonstrating the Agilex Ranger Mini 2.0 successfully navigating to its target, a hair brush. By anticipating that the shortest path is blocked, our method proactively plans a longer, yet feasible, alternative route to circumvent the obstacle.

Similarly, Fig. 15 illustrates a search for a “healthy vegetable” where the initial predicted location is empty. Instead of relying on the outdated map, PredictiveGraphs correctly predicts a new location for the object, allowing the agent to recover and successfully complete the task.

These experiments are also shown in the first section of the video that we include as part of our supplementary material.




Fig. 15. **Qualitative Navigation Result in Real-World - Replanning:** Open-vocabulary query “Could you get me a healthy vegetable?” with query time 29 hours after the last map update. (1) PredictiveGraphs identifies the target object as a carrot. (2) PredictiveGraphs proposes an initial candidate location; however, upon navigating to it, the object is not found. (3) Instead of reverting to the object’s last known location, PredictiveGraphs predicts a new location where the carrot is successfully found, highlighting the robustness of our tempo-spatio-semantic scene graph in dynamic scenes.

REFERENCES

- [1] C. Stachniss and W. Burgard, “Mobile robot mapping and localization in non-static environments,” in *AAAI*, 2005, pp. 1324–1329.
- [2] P. Biber and T. Duckett, “Dynamic maps for long-term operation of mobile service robots,” in *Robotics: Science and Systems*, 2005.
- [3] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Trans. on Robot.*, pp. 1309–1332, 2016.
- [4] L. Carlone, A. Kim, T. Barfoot, D. Cremers, and F. Dellaert, Eds., *SLAM Handbook. From Localization and Mapping to Spatial Intelligence*. Cambridge University Press, 2025.
- [5] L. Schmid, M. Abate, Y. Chang, and L. Carlone, “Khronos: A unified approach for spatio-temporal

Scene: 153 | Area: 122.6 m² | Rooms: 6



Scene: 172 | Area: 102.9 m² | Rooms: 5



Scene: 281 | Area: 104.2 m² | Rooms: 5



Scene: 336 | Area: 115.1 m² | Rooms: 6



Scene: 453 | Area: 126.7 m² | Rooms: 5



Scene: 534 | Area: 139.8 m² | Rooms: 5



Scene: 559 | Area: 136.0 m² | Rooms: 6



Scene: 609 | Area: 126.7 m² | Rooms: 5



Scene: 612 | Area: 149.2 m² | Rooms: 5



Scene: 694 | Area: 130.5 m² | Rooms: 5



Scene: 710 | Area: 108.5 m² | Rooms: 5



Scene: 745 | Area: 118.4 m² | Rooms: 5



Scene: 858 | Area: 128.3 m² | Rooms: 6



Scene: 967 | Area: 102.3 m² | Rooms: 5



Scene: 976 | Area: 137.5 m² | Rooms: 5



Fig. 16. **ProcTHOR environments:** Top-down view of the different 15 ProcTHOR scenes used in the simulation experiments in Sec VI-A2. Each scene has at least 5 rooms and a navigable space between 100 m² and 150 m².

- metric-semantic slam in dynamic environments,” in *International Journal of Robotics Research*, 2024.
- [6] Z. Lu, J. Ye, and J. Leonard, “3dgs-cd: 3d gaussian splatting-based change detection for physical object rearrangement,” *IEEE Robot. and Automation Lett.*, 2025.
 - [7] L. Schmid *et al.*, “Panoptic multi-tsdfs: a flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2022, pp. 8018–8024.
 - [8] A. Anwar, J. Welsh, J. Biswas, S. Pouya, and Y. Chang, “Remembr: Building and reasoning over long-horizon spatio-temporal memory for robot navigation,” 2025, pp. 2838–2845.
 - [9] Q. Xie, S. Y. Min, P. Ji *et al.*, “Embodied-rag: General non-parametric embodied memory for retrieval and generation,” *Preprint*, 2024.
 - [10] D. Rosen, J. Mason, and J. Leonard, “Towards lifelong feature-based mapping in semi-static environments,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2016, pp. 1063–1070.
 - [11] F. Nobre, C. Heckman, P. Ozog, R. W. Wolcott, and J. M. Walls, “Online probabilistic change detection in feature-based maps,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2018, pp. 3661–3668.
 - [12] M. Saavedra-Ruiz, S. Nashed, C. Gauthier, and L. Paull, “Perpetua: Multi-hypothesis persistence modeling for semi-static environments,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2025.
 - [13] T. Krajník, J. P. Fentanes, J. M. Santos, and T. Duckett, “Fremen: Frequency map enhancement for long-term mobile robot autonomy in changing environments,” *IEEE Trans. on Robot.*, pp. 964–977, 2017.
 - [14] Y. Wang, Y. Fan, J. Wang, and W. Chen, “Long-term navigation for autonomous robots based on spatio-temporal map prediction,” *Robotics and Autonomous Systems*, p. 104724, 2024.
 - [15] G. Comanici *et al.*, “Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities,” *Preprint*, 2025.
 - [16] Q. Gu *et al.*, “Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2024, pp. 5021–5028.
 - [17] V. Yugay, T. Kersten, L. Carlone, T. Gevers, M. R. Oswald, and L. Schmid, “Gaussian mapping for evolving scenes,” *Preprint*, 2025.
 - [18] J. Qian *et al.*, “Pov-slam: Probabilistic object-level variational slam,” in *Robotics: Science and Systems*, 2023.
 - [19] T. Behrens, R. Zurbrugg, M. Pollefeys, Z. Bauer, and H. Blum, “Lost & found: Tracking changes from ego-centric observations in 3d dynamic scene graphs,” *IEEE Robot. and Automation Lett.*, pp. 3739–3746, 2025.
 - [20] V. Guizilini, R. Senanayake, and F. Ramos, “Dynamic hilbert maps: Real-time occupancy predictions in changing environments,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2019, pp. 4091–4097.
 - [21] T. Deng, H. Xie, J. Wang, and W. Chen, “Long-term visual simultaneous localization and mapping: Using a bayesian persistence filter-based global map prediction,” *RAM*, pp. 36–49, 2023.
 - [22] W. Ali, Y. Cai, P. Jensfelt, and T.-M. Nguyen, “Proberlilo: Probabilistic persistency modeling for life-long mapping,” pp. 2530–2537, 2026.
 - [23] A. Kurenkov *et al.*, “Modeling dynamic environments with scene graph memory,” in *International Conference on Machine Learning*, 2023, pp. 17 976–17 993.
 - [24] S. Looper, J. Rodriguez-Puigvert, R. Siegwart, C. Cadena, and L. Schmid, “3d vsg: Long-term semantic scene change prediction through 3d variable scene graphs,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2023, pp. 8179–8186.
 - [25] M. Patel and S. Chernova, “Proactive robot assistance via spatio-temporal object modeling,” *Preprint*, 2022.
 - [26] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, “Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation,” in *VLMNM Workshop at ICRA*, 2024.
 - [27] D. Maggio *et al.*, “Clio: Real-time task-driven open-set 3d scene graphs,” 2024.
 - [28] B. Bogenberger *et al.*, “Where did i leave my glasses? open-vocabulary semantic exploration in real-world semi-static environments,” *PrePrint*, 2025.
 - [29] P. Liu *et al.*, “Dynamem: Online dynamic spatio-semantic memory for open world mobile manipulation,” 2025, pp. 13 346–13 355.
 - [30] P. Lewis *et al.*, “Retrieval-augmented generation for knowledge-intensive nlp tasks,” *NeurIPS*, pp. 9459–9474, 2020.
 - [31] M. F. Ginting *et al.*, “Enter the mind palace: Reasoning and planning for long-term active embodied question answering,” in *CoRL*, 2025.
 - [32] M. Yuan, H. Zhang, M. Mohammadi, R. Li, J. Shan, and S. L. Waslander, “Star: Scalable task-conditioned retrieval for long-horizon multi-modal robot memory,” *IEEE Robot. and Automation Lett.*, pp. 5994–6001, 2026.
 - [33] J. Jiang, Y. Zhu, Z. Wu, and J. Song, “Dualmap: Online open-vocabulary semantic mapping for natural language navigation in dynamic changing scenes,” *IEEE Robot. and Automation Lett.*, pp. 12 612–12 619, 2025.
 - [34] Y. Tang *et al.*, “Openin: Open-vocabulary instance-oriented navigation in dynamic domestic environments,” *IEEE Robot. and Automation Lett.*, pp. 9256–9263, 2025.
 - [35] K. P. Murphy, *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.
 - [36] C. Wang, dong wang, X. Li, D. Wang, and H. Liu, “Dynamic scene generation for embodied navigation benchmark,” in *RSS Workshop: Data Generation for Robotics*, 2024.
 - [37] A. Radford, J. W. Kim, C. Hallacy *et al.*, “Learning transferable visual models from natural language supervision,” in *ICML*, 2021, pp. 8748–8763.
 - [38] M. Deitke *et al.*, “Proctor: Large-scale embodied ai using procedural generation,” pp. 5982–5994, 2022.
 - [39] M. Savva *et al.*, “Habitat: A platform for embodied ai research,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019,

pp. 9339–9347.

- [40] M. Labbé and F. Michaud, “Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation,” *Journal of field robotics*, pp. 416–446, 2019.
- [41] H. Akaike, “A new look at the statistical model identification,” in *Selected Papers of Hirotugu Akaike*. Springer, 1974, pp. 215–222.
- [42] N. Carion *et al.*, “Sam 3: Segment anything with concepts,” *Preprint*, 2025.