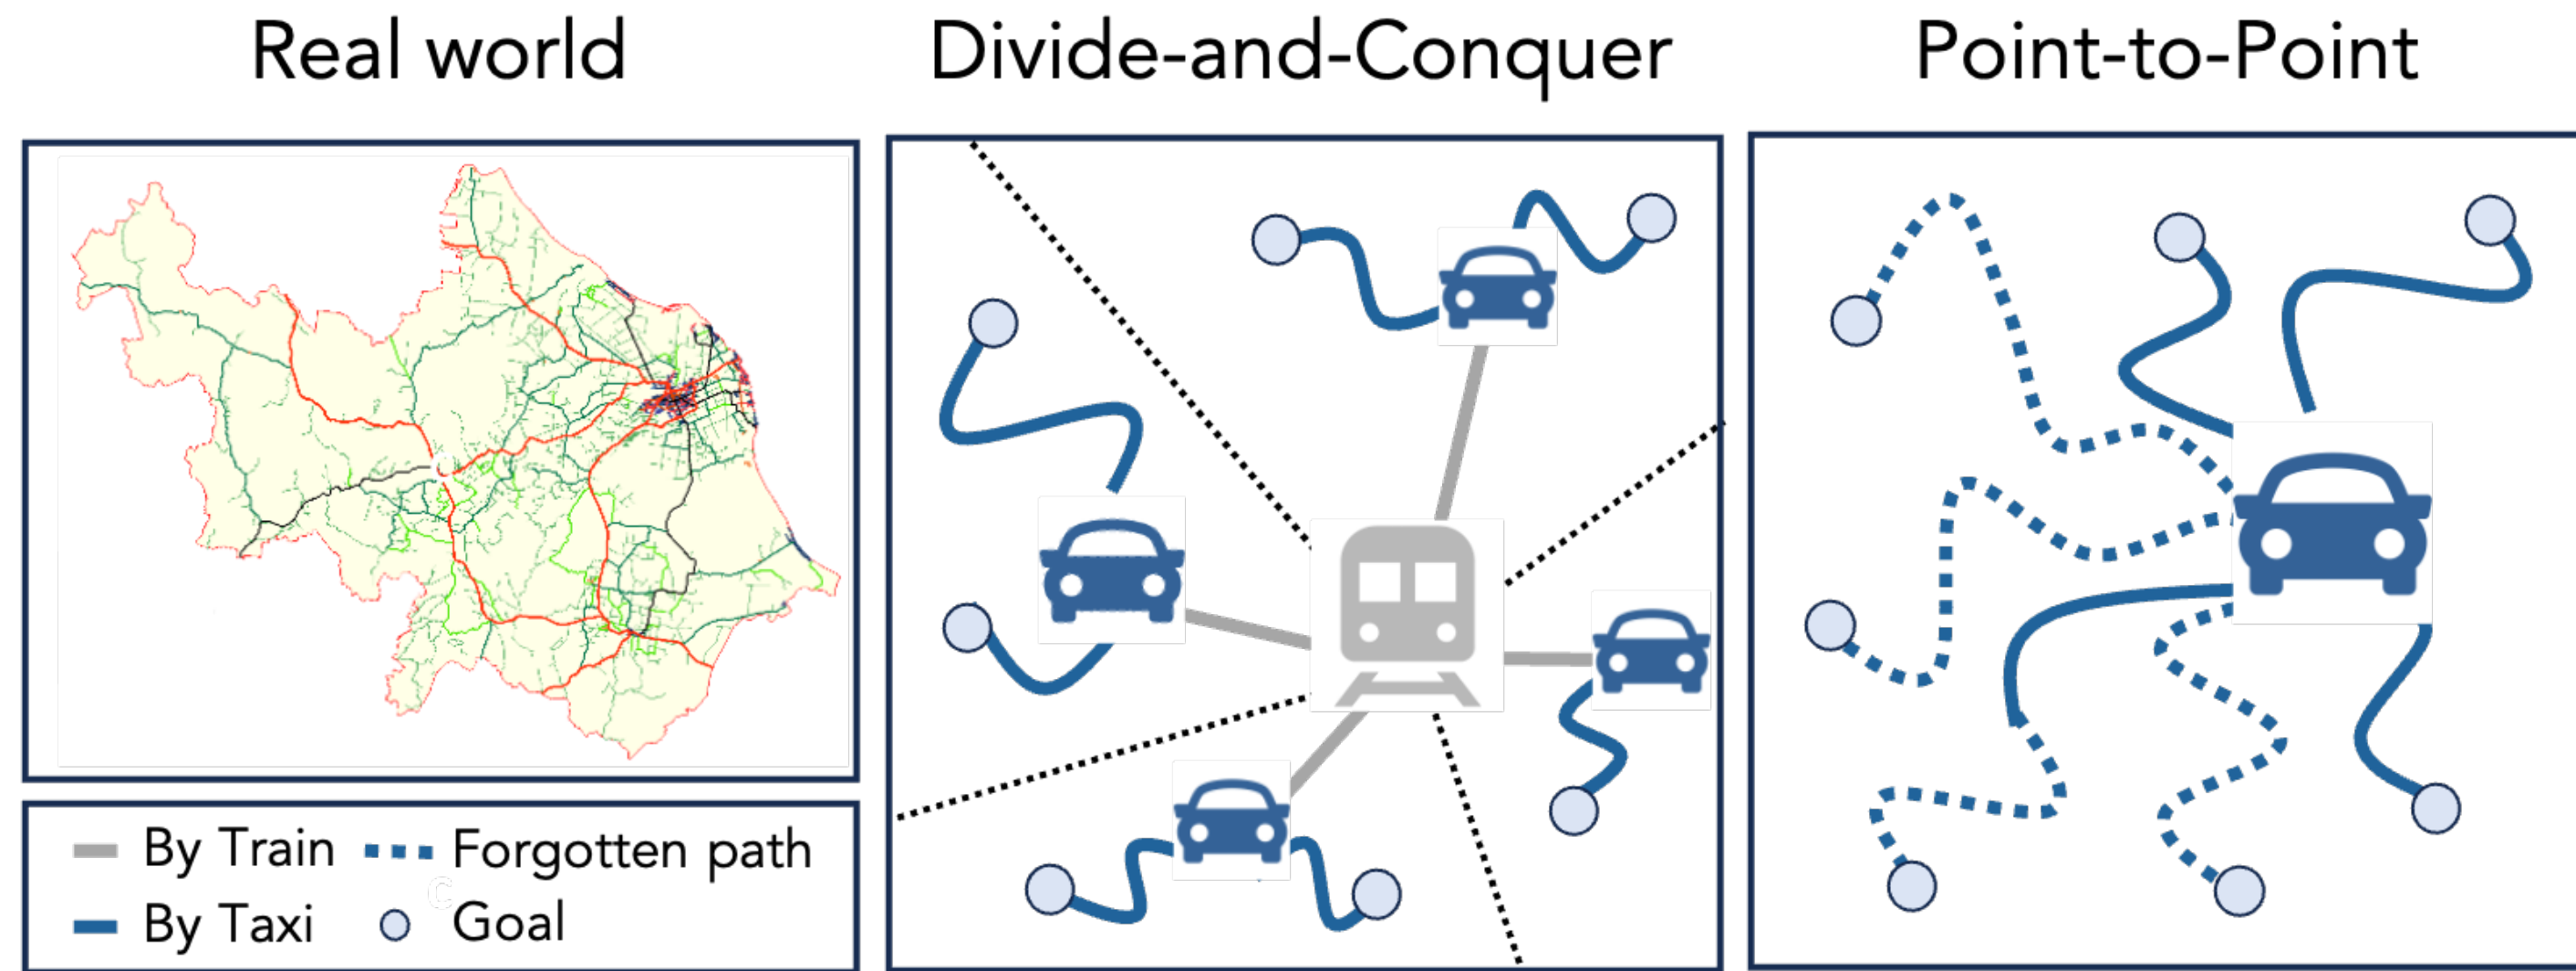




Generalist Agents

A crucial capability of generalist agents, such as humans, is to explore environments and acquire the abilities needed to achieve various goals, continuously and in an open-ended way.

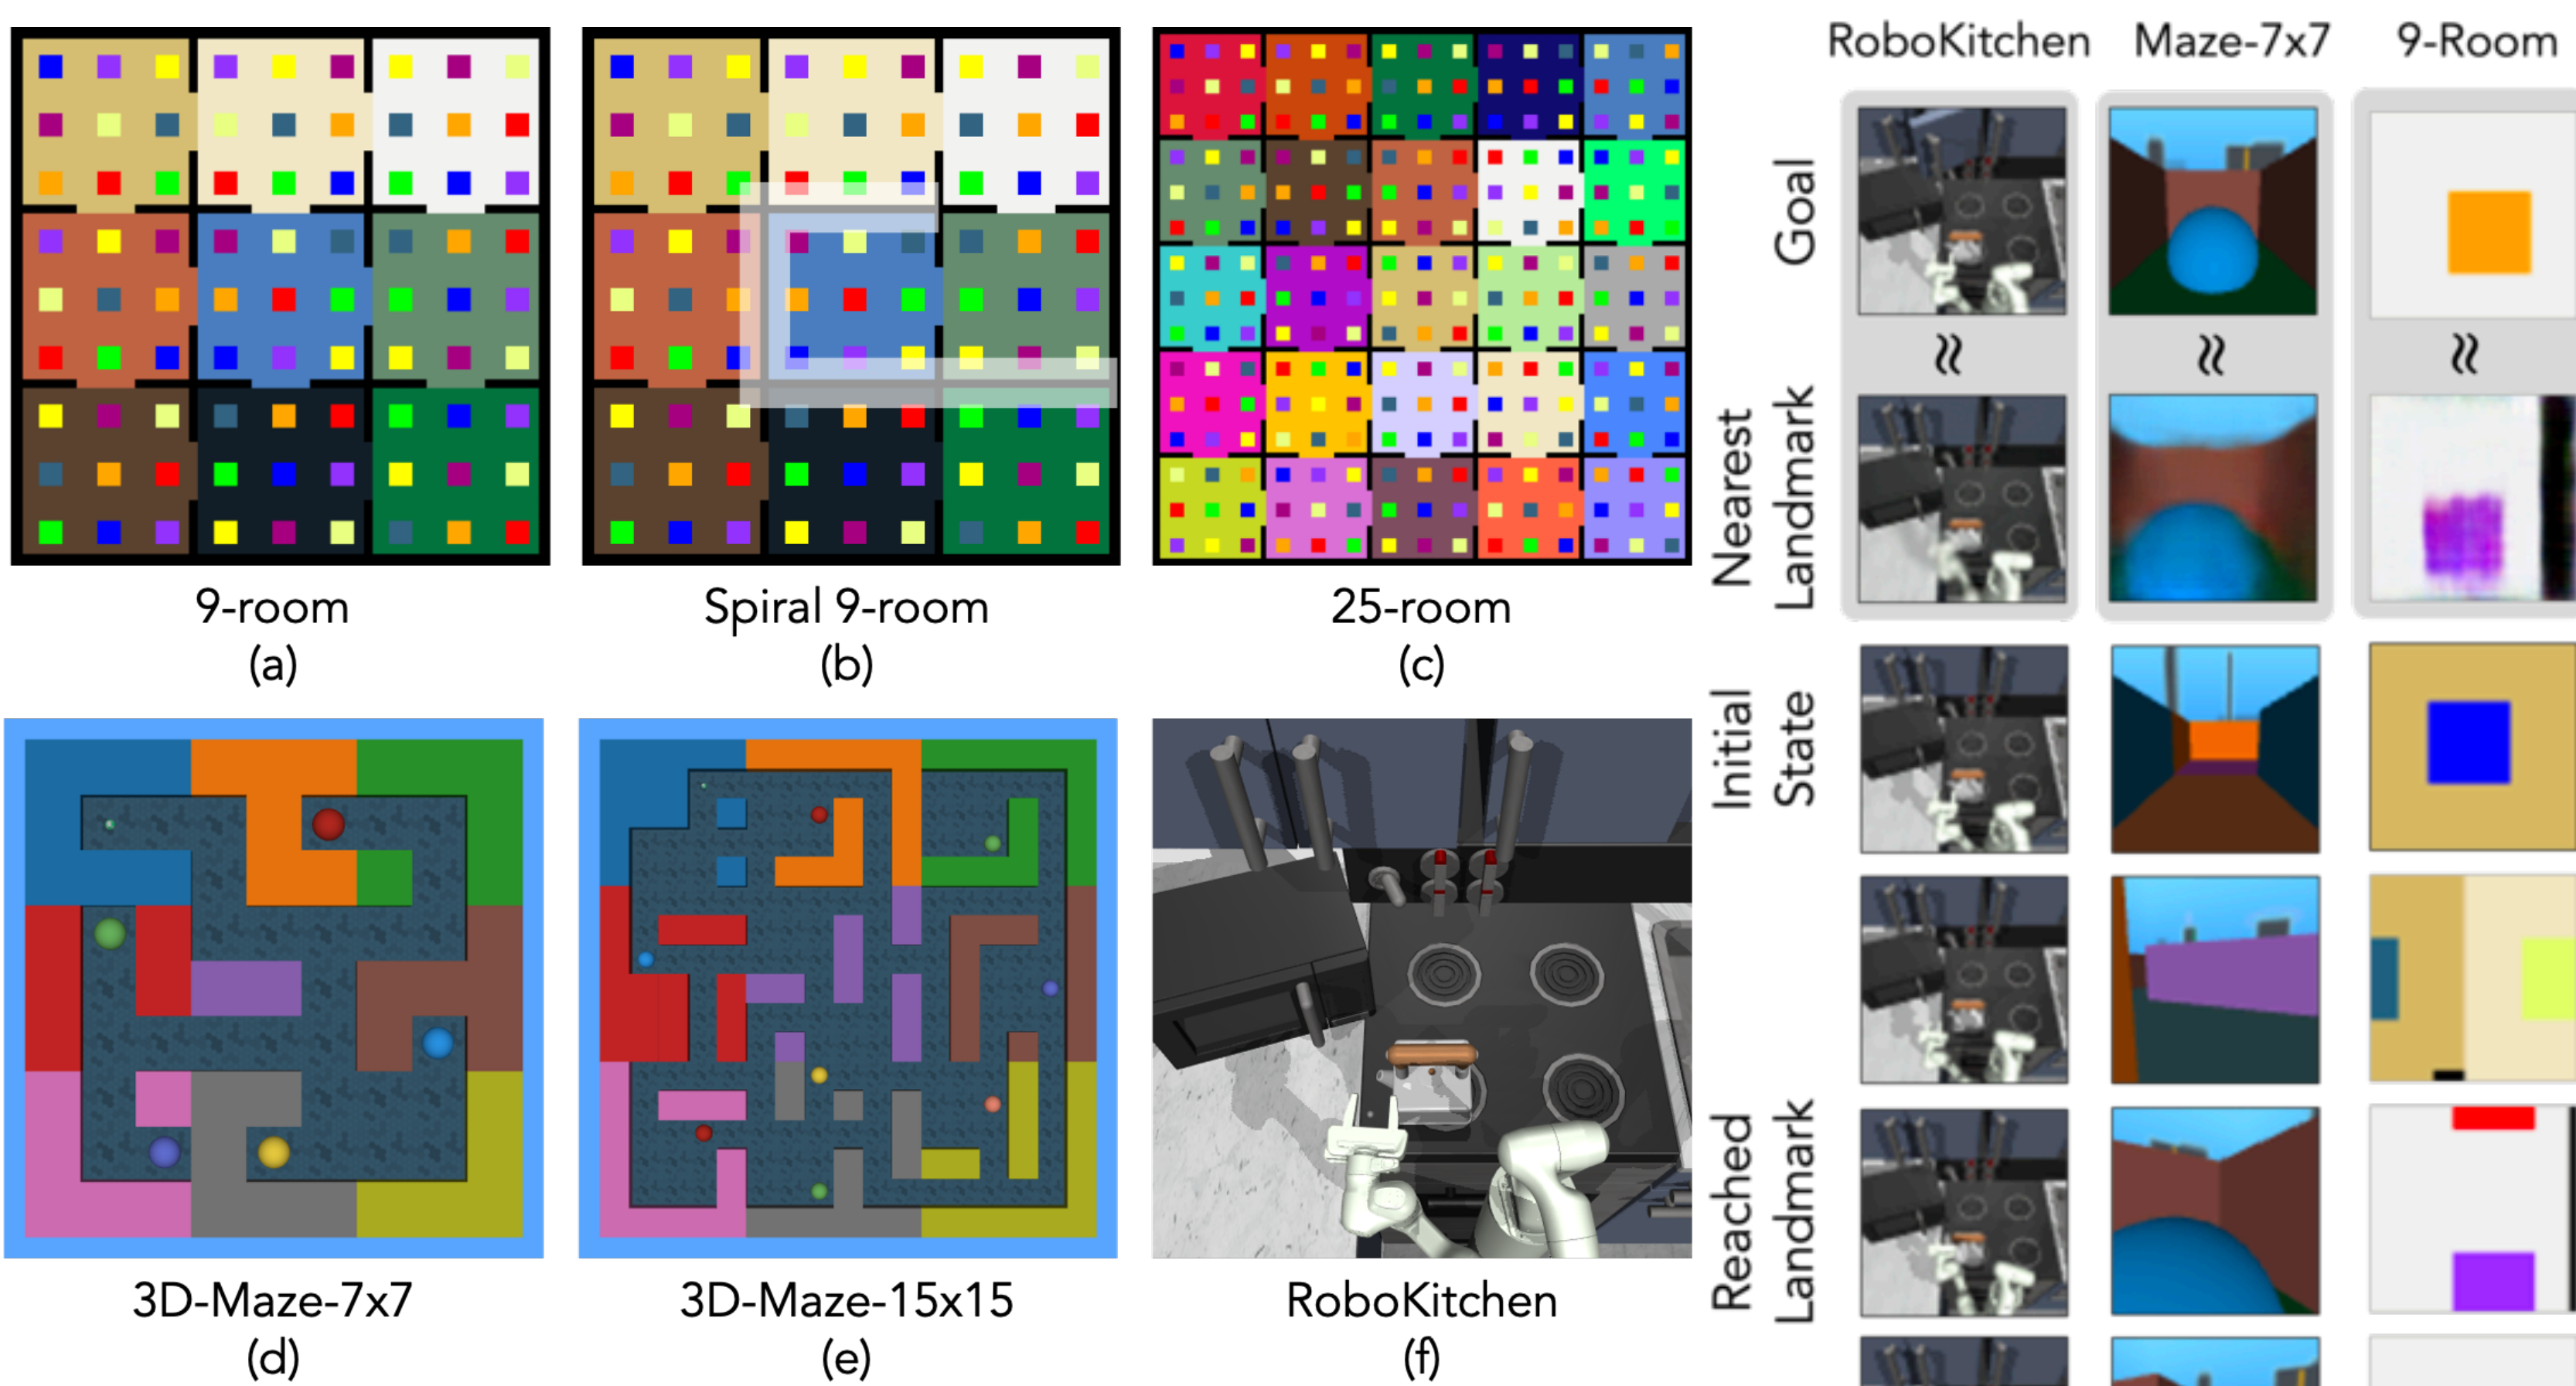
We raise the question of whether and how a model-based generalist RL agent can “dream better” in a more structured and strategic way.



We propose a strategic model-based generalist agent, Dr. Strategy (short for “Dream Strategy”)

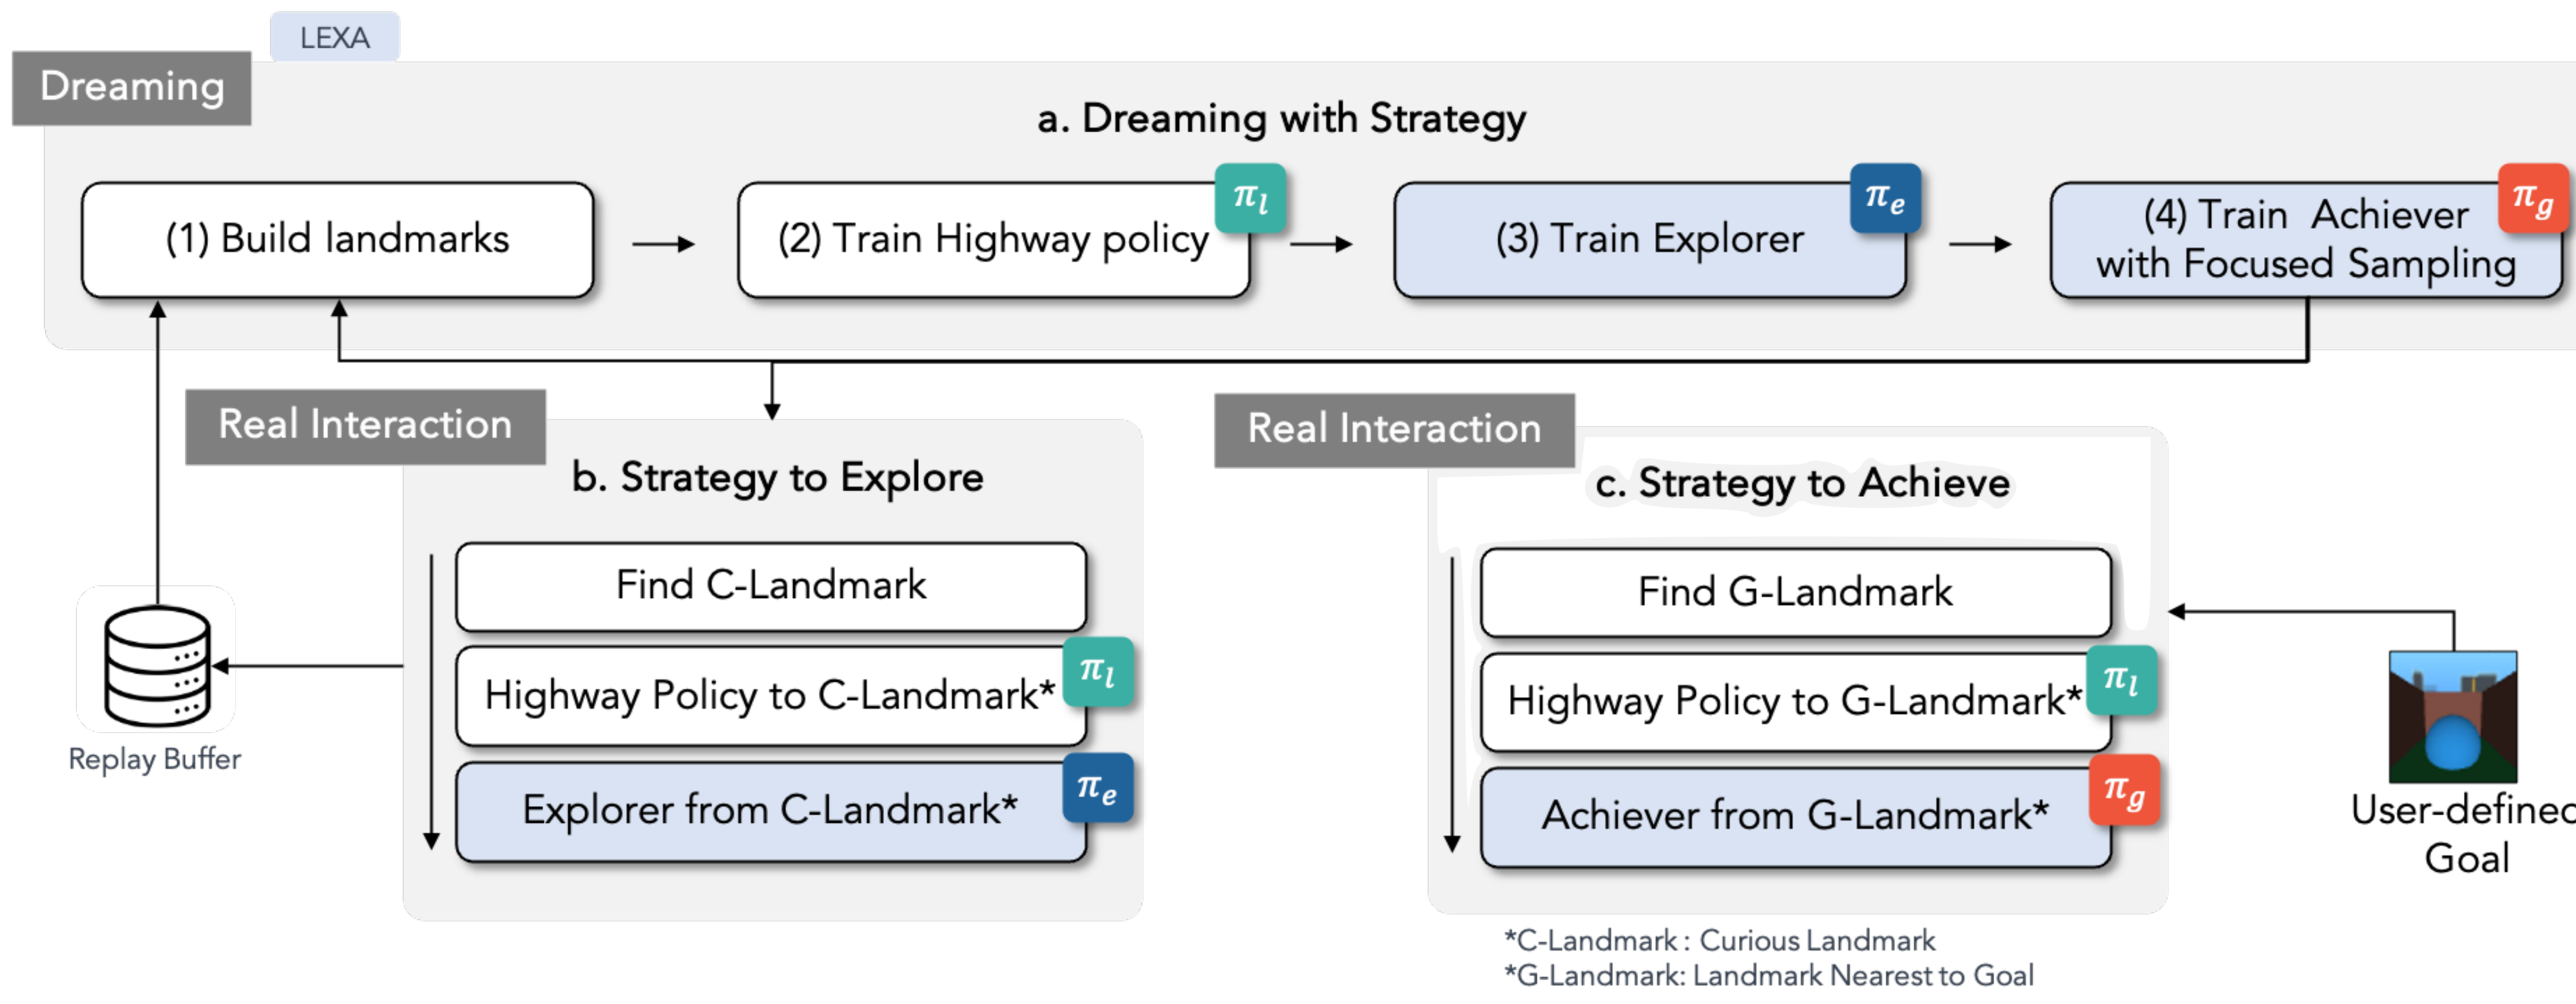
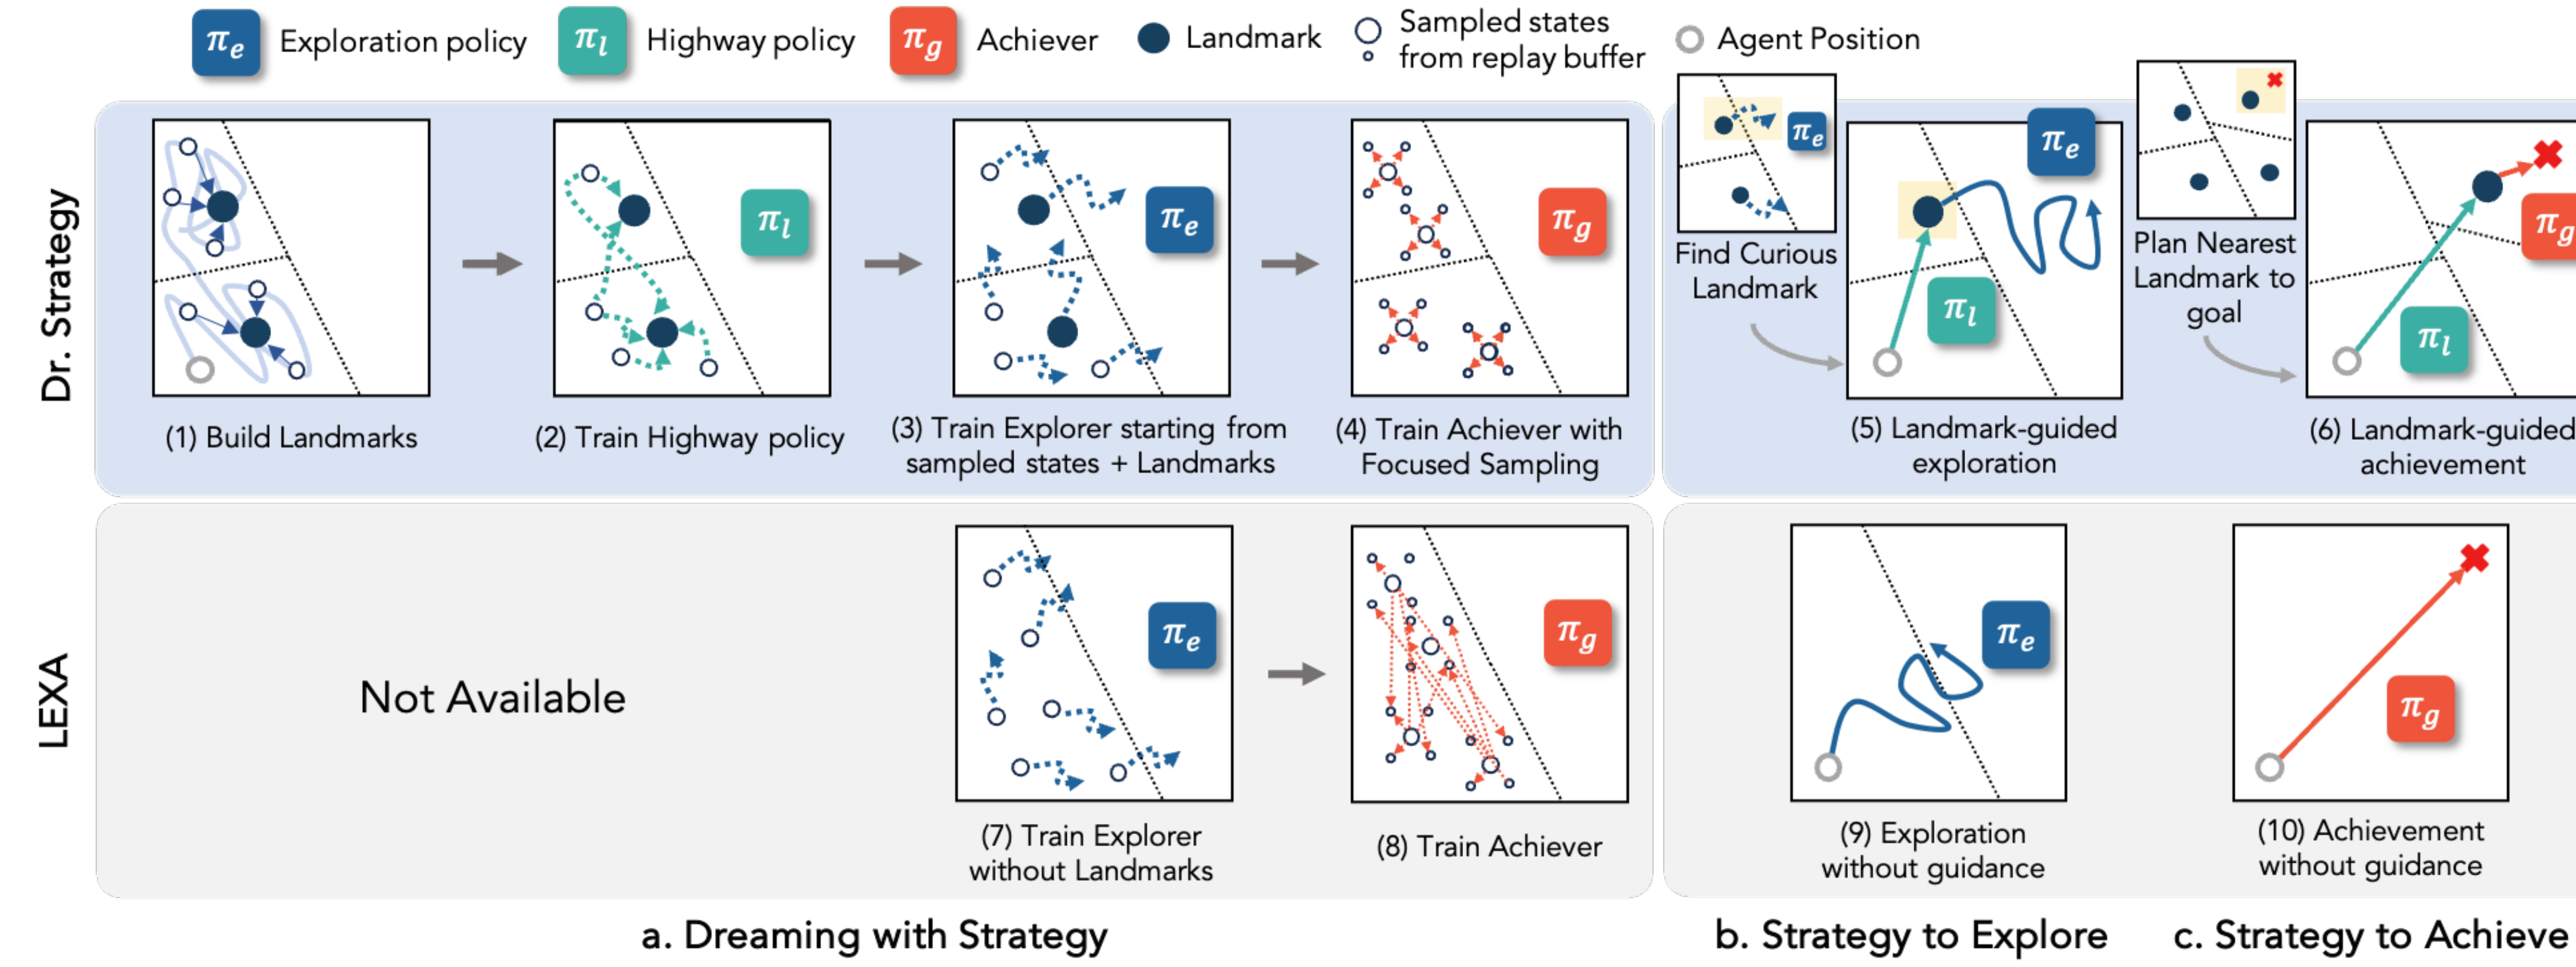
Our key idea is that a *divide-and-conquer* approach leveraging the structure of latent landmarks can enhance the efficiency of dreaming in MBRL and promote better exploration and achievement quality of a generalist agent.

Environments



Three types of visual-based environments:
 (a-c) 2D POMDP navigation (partial view)
 (d-e) 3D-Maze Navigation (ego-centric)
 (f) a Franka Panda arm manipulation (third-person view)

Dr. Strategy



Dreaming with Strategy

- Train the world model (DreamerV2)

$$\text{Dynamics: } \hat{s}_t \sim \text{dyn}_\theta(s_{t-1}, a_{t-1}) \quad \text{Encoder: } e_t = \text{enc}_\theta(x_t)$$

$$\text{Representation: } s_t \sim \text{repr}_\theta(s_{t-1}, a_{t-1}, x_t) \quad \text{Decoder: } \hat{x}_t \sim \text{dec}_\theta(s_t)$$

- Learn the latent landmarks from the replay buffer using VQ-VAE

- Train three policies in imagination (Strategic Dreaming)

$$1) \text{ Highway policy: } r_l(s_t, l) = -\|\text{dec}_\phi(l) - s_t\|_2^2 + \sum_{i=1}^K \log\|s_t - s_i^{\text{K-NN}}\|_2$$

$$2) \text{ Explorer: } r_e(s_t) \propto \text{Var}(\{f(e_{t+1} | s_t, a_t, \theta_k) | k \in [1, \dots, N]\})$$

$$3) \text{ Achiever: } r_g(\hat{e}_t, e_g) = -\text{tdp}(\hat{e}_t, e_g)$$

- Calculate the landmarks' curiosity: $C_i = \mathbb{E}_{\tau_i}[r_e(s_t^{(i)})]$, $\tau_i = \{s_0^{(i)}, s_1^{(i)}, \dots, s_H^{(i)}\}$

Results

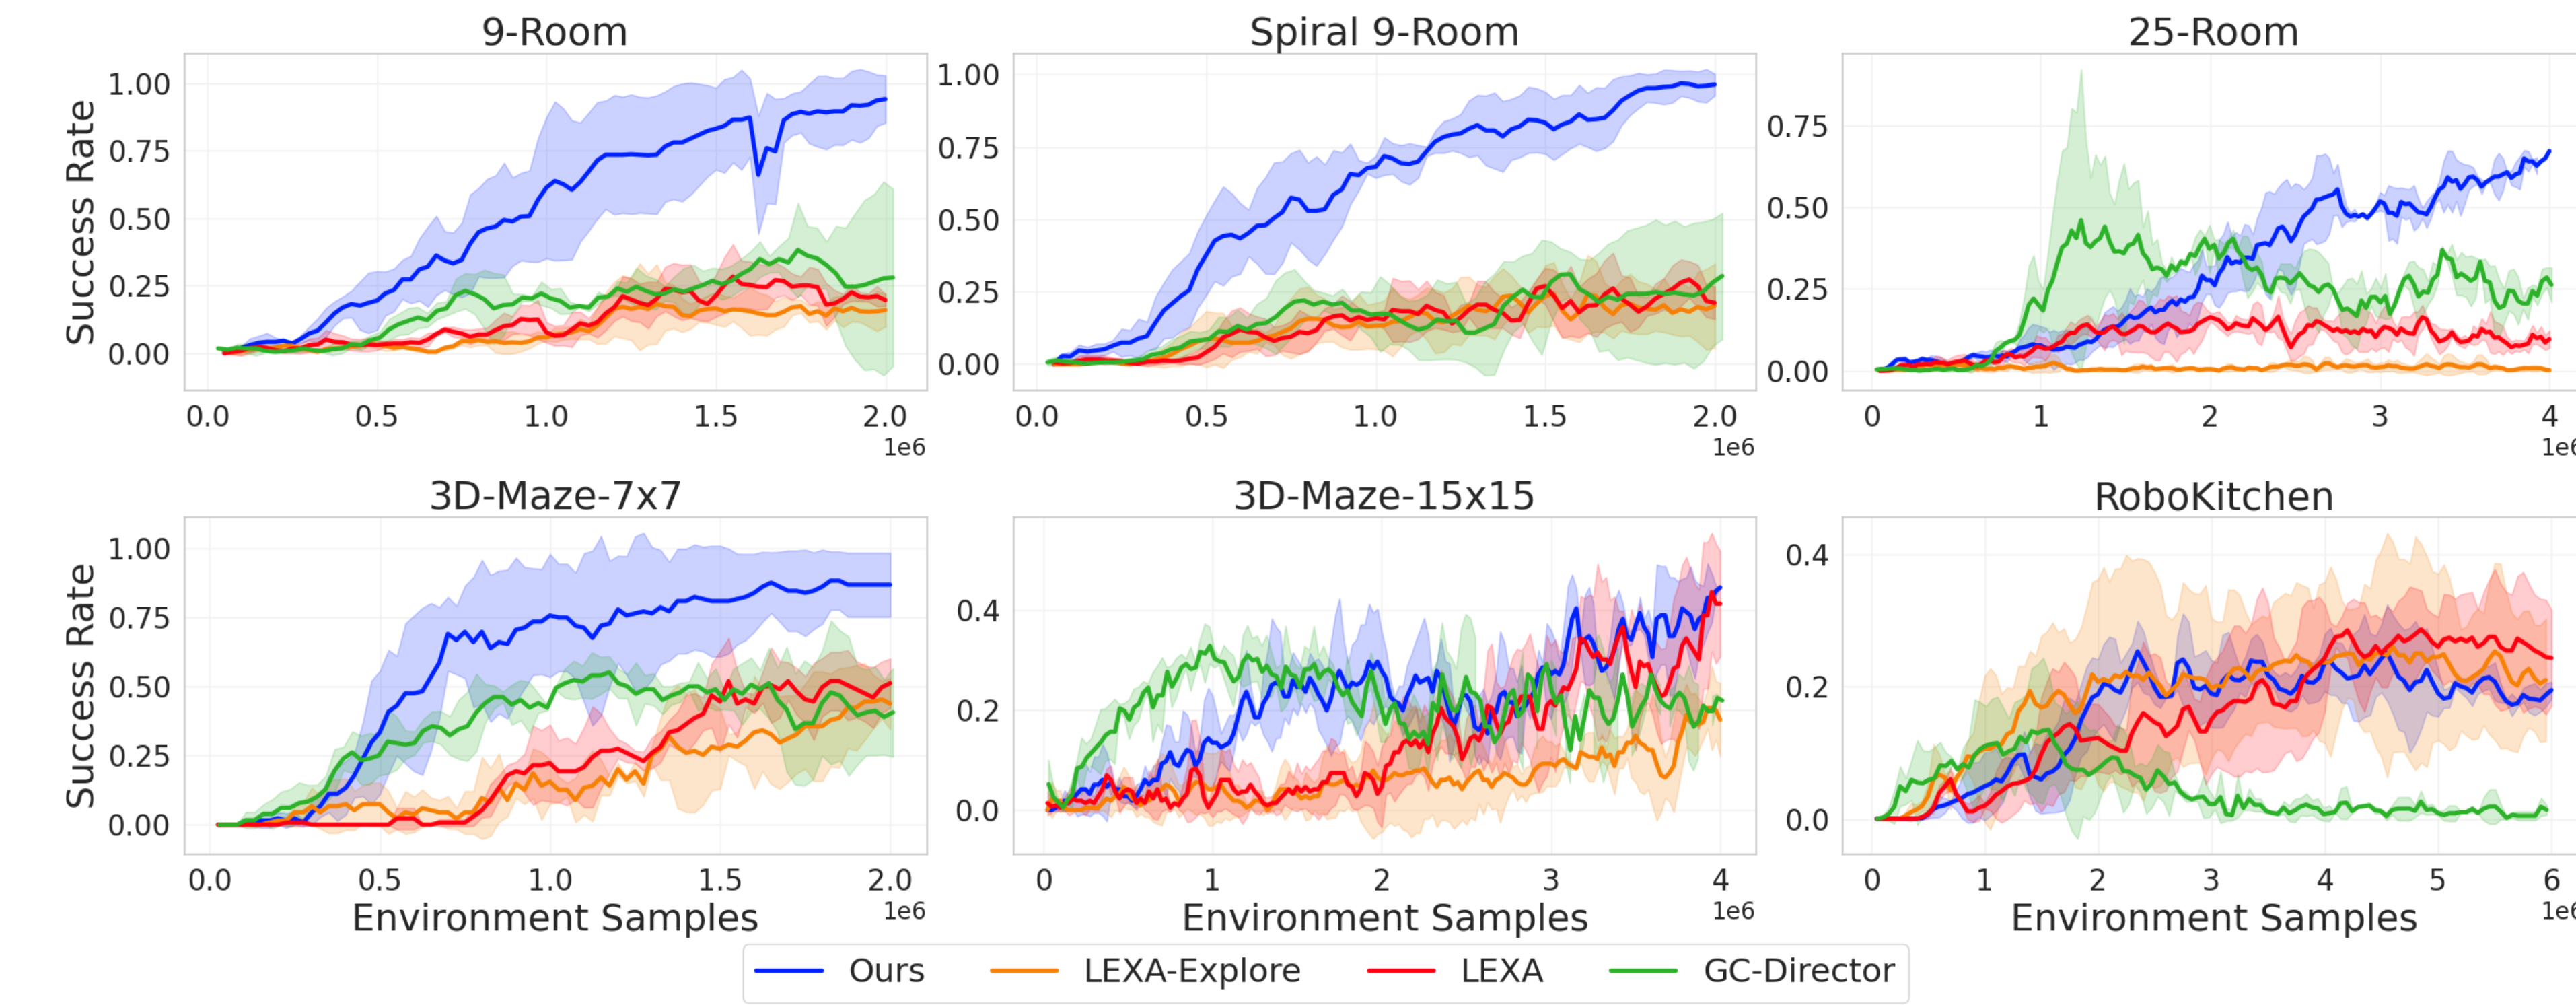


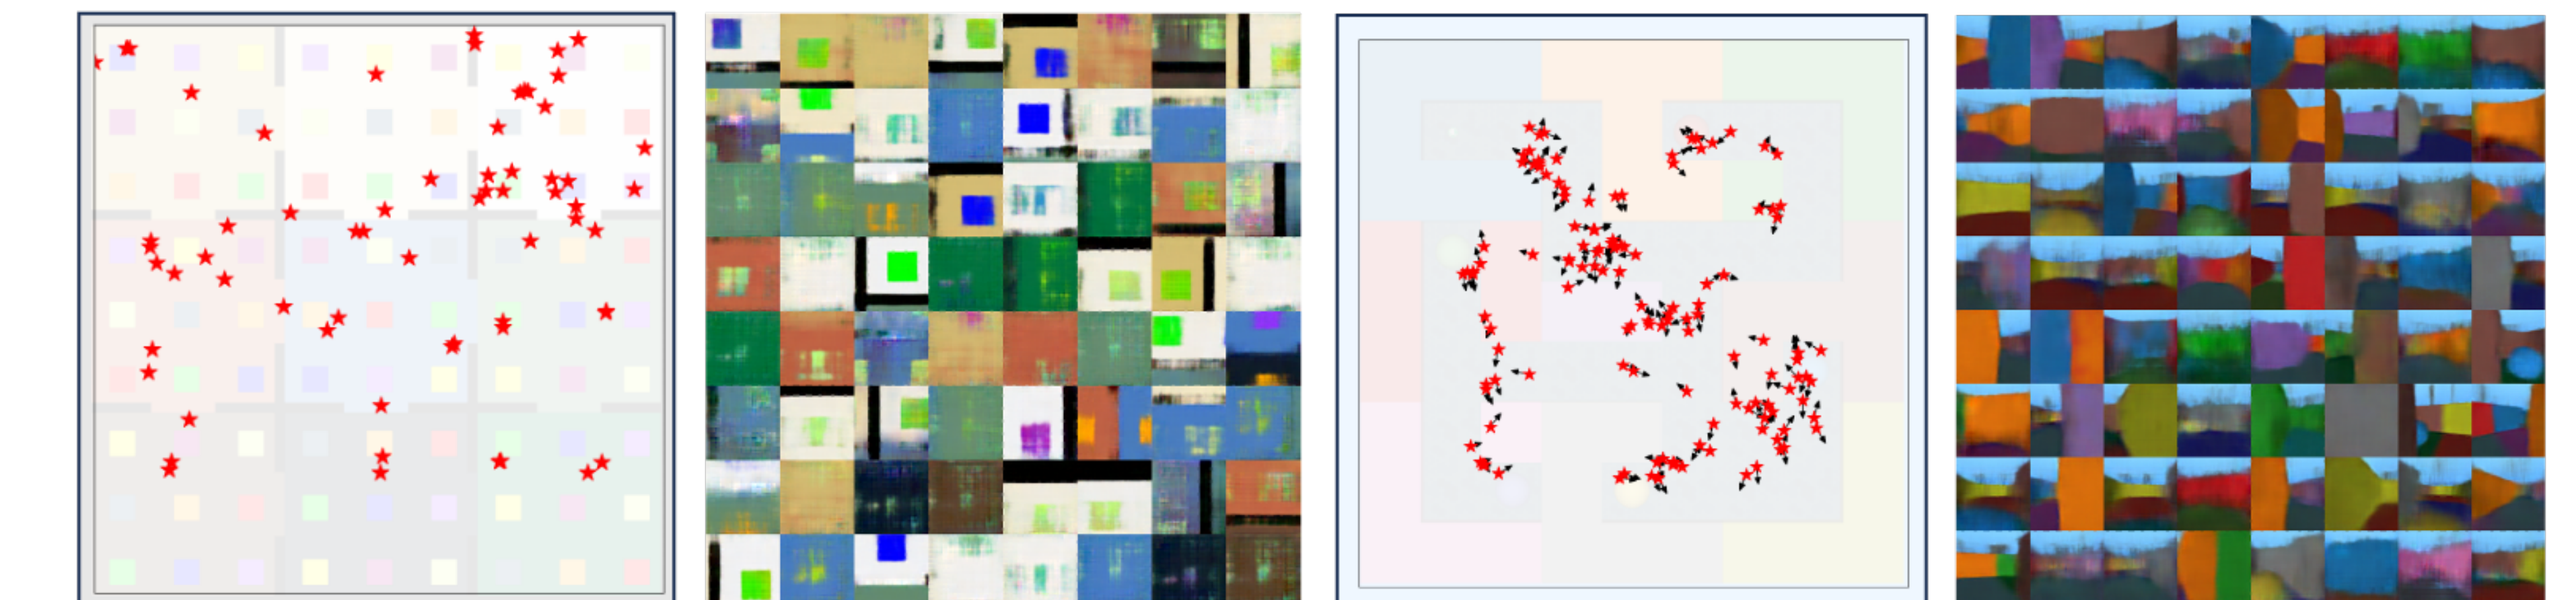
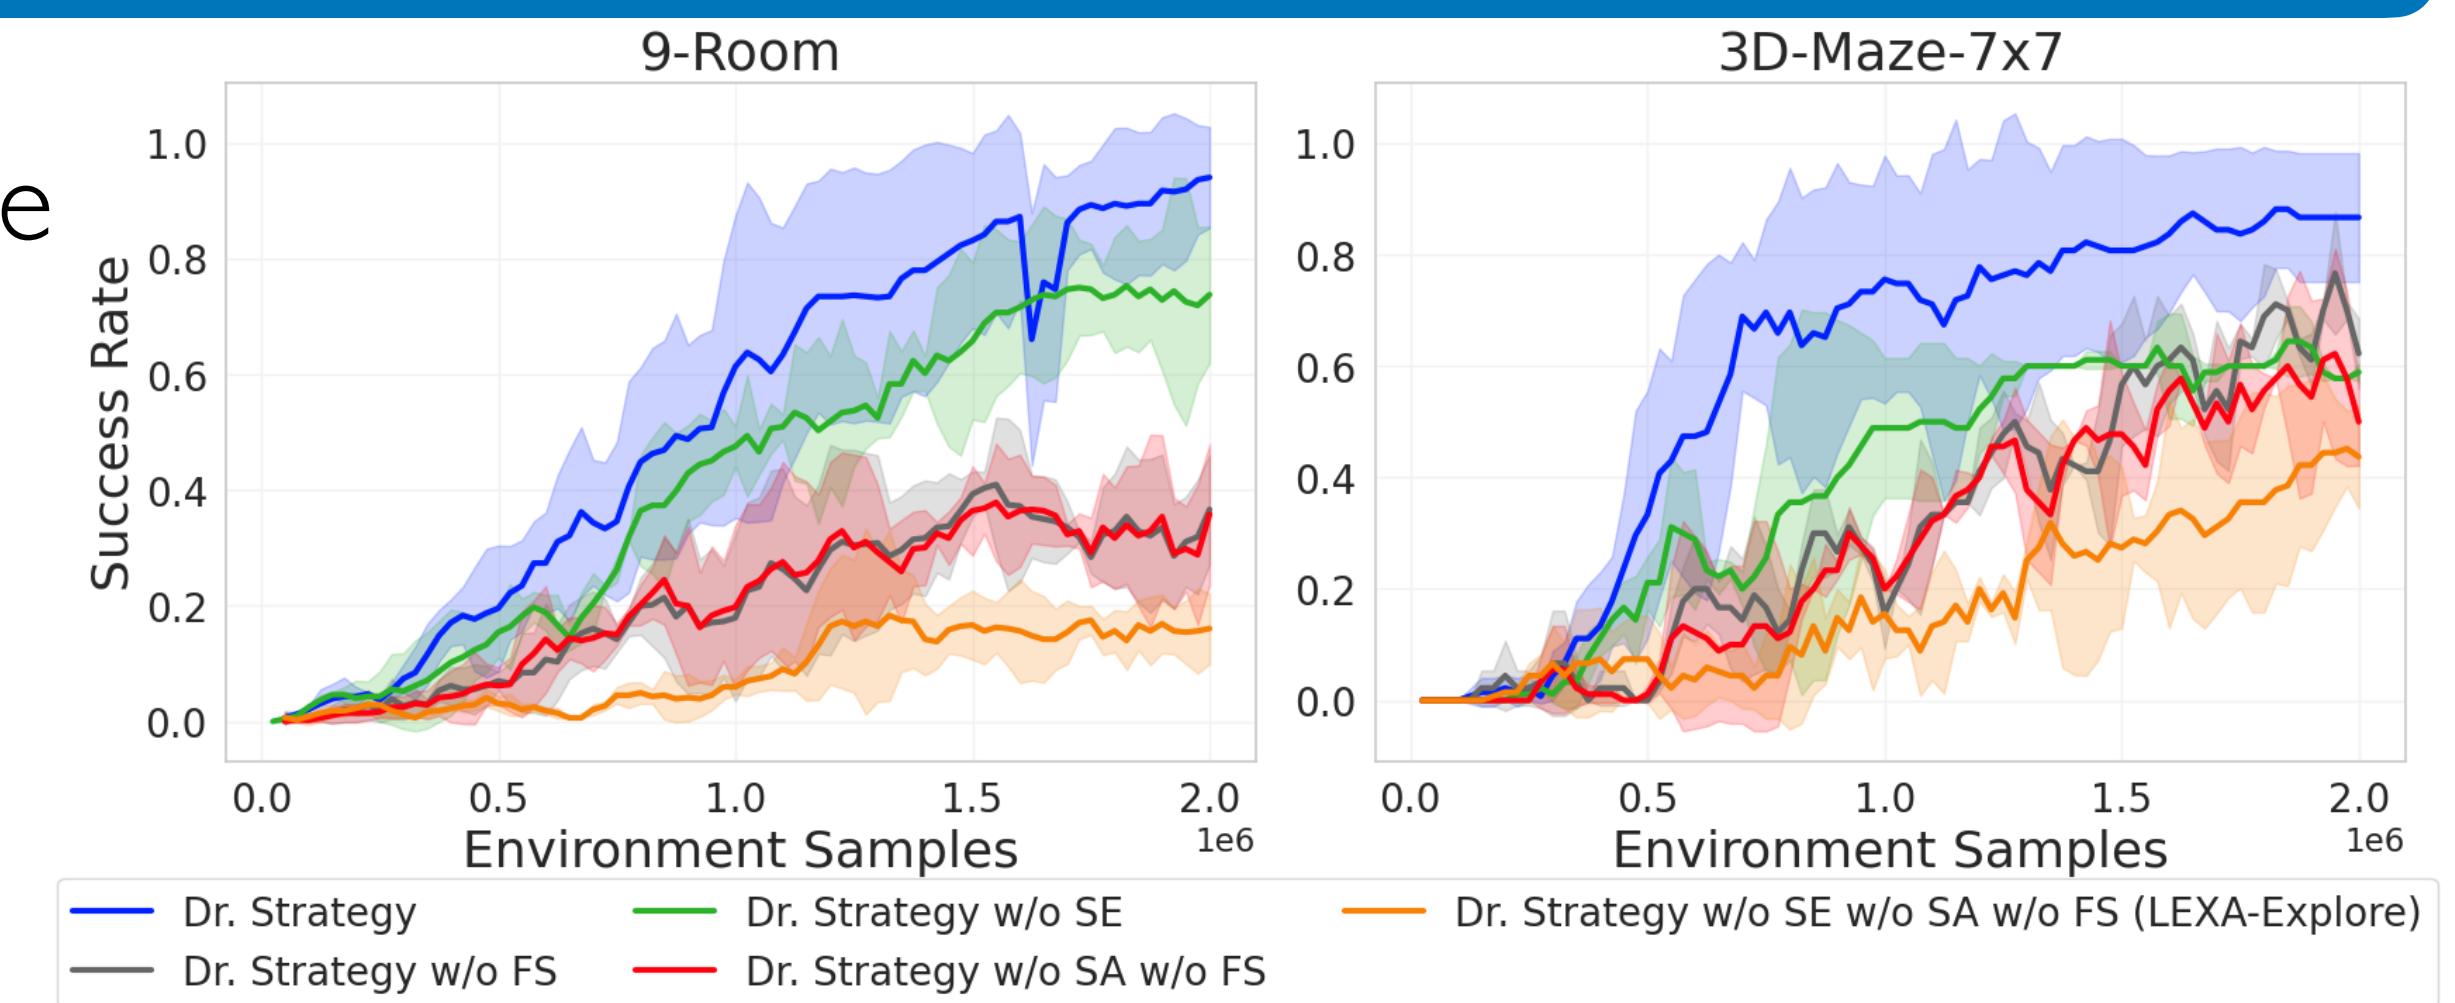
Table 1. Final Success Rate

Method	9-Room	Spiral 9-Room	25-Room	Maze-7x7	Maze-25x25	RoboKitchen
LEXA	19.75%	21.19%	9.62%	51.11%	41.20%	24.30%
LEXA-Explore	16.04%	20.16%	0.14%	43.70%	18.05%	20.07%
GC-Director	28.08%	30.45%	27.11%	40.55%	21.87%	1.45%
Dr. Strategy (Ours)	94.03%	96.50%	67.11%	86.66%	44.44%	19.44%

Ablation Studies

We investigate the influence of three components of strategic dreaming:

Strategy to Explore (SE),
 Strategy to Achieve (SA),
 Focused Sampling (FS)



Reconstructed Landmarks for 9-Room and 3D-Maze 7x7

Conclusion

- We propose Dr. Strategy, a novel model-based task generalist agent.
- Dr. Strategy outperformed prior pixel-based MBRL agents in various visually complex and partially observable navigation tasks, while also showing comparative results in robot manipulation tasks.