VIBE-GAS: Variational Informed Belief Estimation for Graph Active Search

Online Planning for Active Search on Partially Observed Graphs

Teim Giesen

Supervised by: **Prof. Dr. F.A. (Frans) Oliehoek** & **Andreea Minculescu**

1. Graph Active Search: Motivation

- **Graph Active Search:** sequentially querying a pool of candidates (nodes in a graph) to discover rare targets as quickly as possible [1].
- **Applications:** drug discovery, fraud detection, material sciences, identifying oil spills, healthcare, and robotics.
- **Current Work:** limited by the assumption of a known graph topology, reliance on deterministic feedback, and myopic planning.
- **Using a POMDP:** all of the above can be avoided using a Partially Observable Markov Decision Process (POMDP) [2] to enable principled, non-myopic planning under uncertainty.

  Here, states encode graph topology and node labels, while querying a node yields a noisy observation of its label and local neighbourhood. Belief is maintained over all possible graphs.
- **How to Encode Belief During Planning?** Standard online algorithms such as Partially Observable Monte Carlo Planning (POMCP) [3] use a particle filter to maintain belief. **Yet this struggles in graph domains**.

2. Where Particle Filters Fall Short

- **Curse of Dimensionality:** Particles struggle to accurately reflect high dimensional spaces. For example with just 20 nodes; 2^{210} graphs exist.

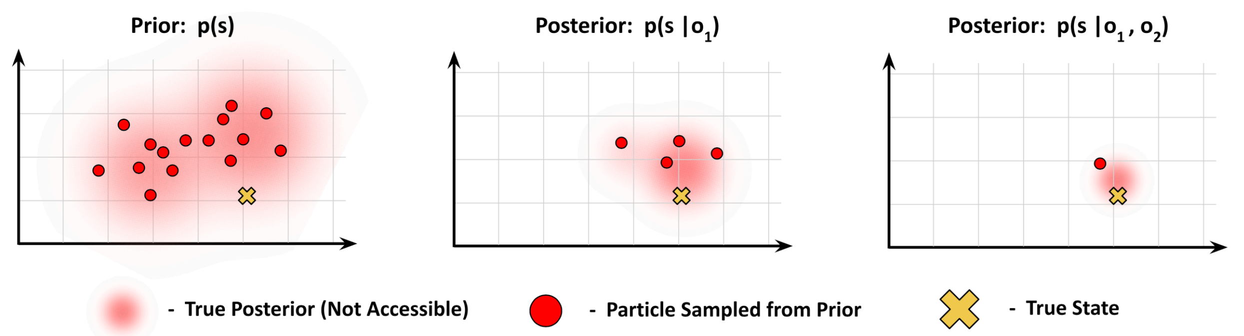

Figure 1: Particle degeneracy visualisation with shrinking posterior.

- **Particle Degeneracy:** It is exponentially unlikely that the true state is captured with random samples. Thus, the filter degenerates to a single (highly unlikely) particle; and Effective Sample Size (ESS) plummets [4].

3. Using Variational Inference

- **Core Idea:** teach a Variational Graph Auto-Encoder (VGAE) [5] to generate particles conditioned on the observation history.

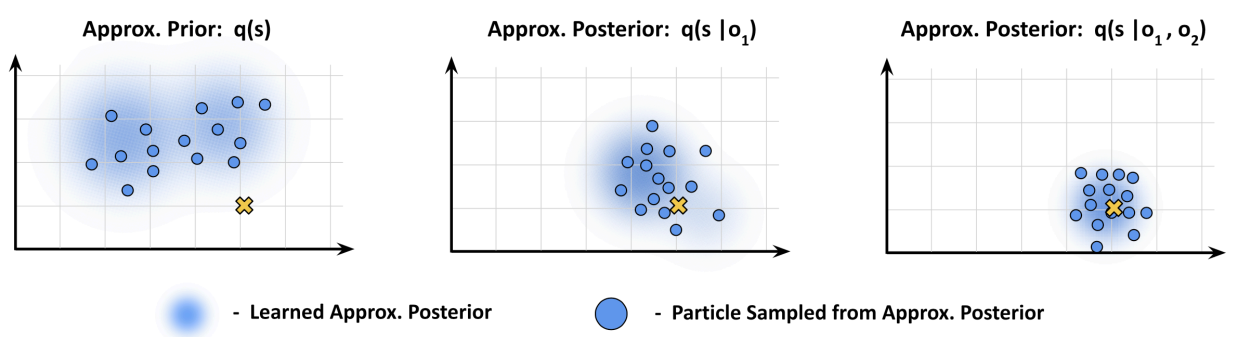

Figure 2: Using an approximate posterior to generate particles.

- **Formally:** Approximate true posterior $p(s \mid h_t)$ with approximate $q(s \mid h_t)$. Draw i.i.d samples from $q(s \mid h_t)$ to act as particles for POMCP planning.
- **Pipeline**: Partial observations → latent space → sample → decode.
- **Training:** The VGAE is trained offline to maximise the ELBO on fully observable graphs with a simulated sequence of observations.
- **Benefits:** No particle degeneracy, native support for node features, learnable bias, inherent diversity within belief, no prior required.

3. VIBE-GAS

- **Incorporating Variational Inference into Active Search:** at each time-step, the agent's action-observation history h_t is ingested by the VGAE. The VGAE synthesises a conditioned fresh particle belief to be exploited by the POMCP planner.
- This yields: **Variational Informed Belief Estimation for Graph Active Search (VIBE-GAS).**

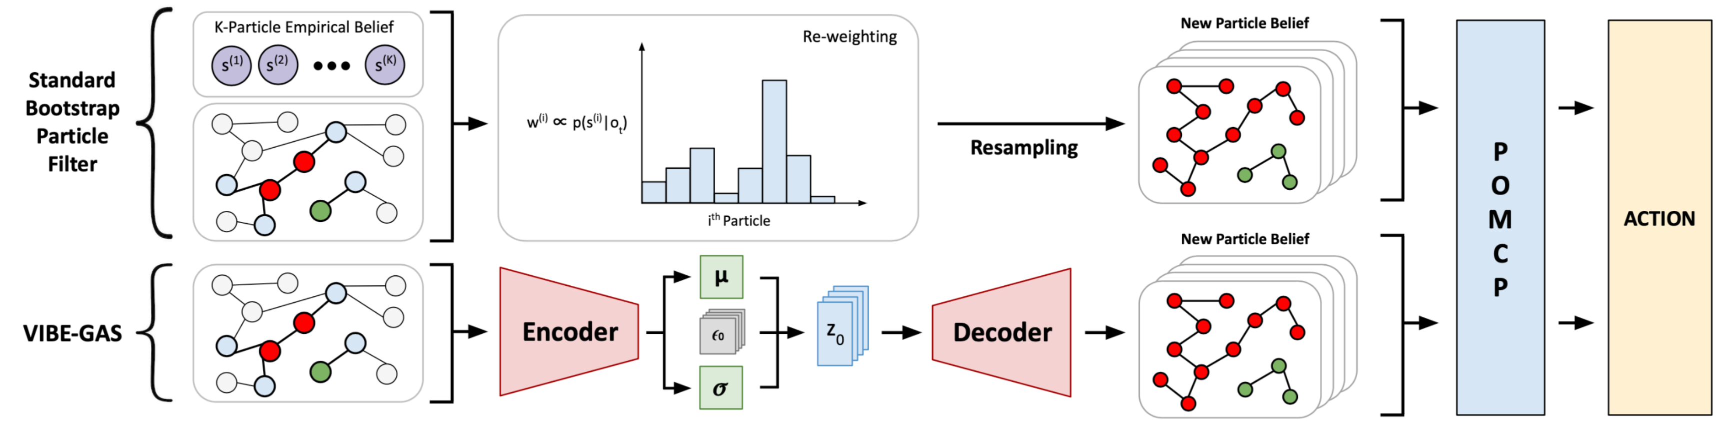

Figure 3: VIBE-GAS vs. Standard Bootstrap Particle Filter Planning

4. Empirical Validation

4.1 Experimental Setup

Dataset: Stochastic Block Model (SBM), Random Geometric (RGG), and Core-Periphery (CP) graphs.

Trials: 20 seeds, 3 runs per seed per algorithm.

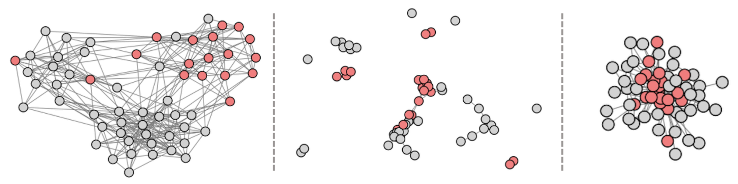

Figure 4: Graphs: (Left) SBM, (Middle) RGG, and (Right) CP

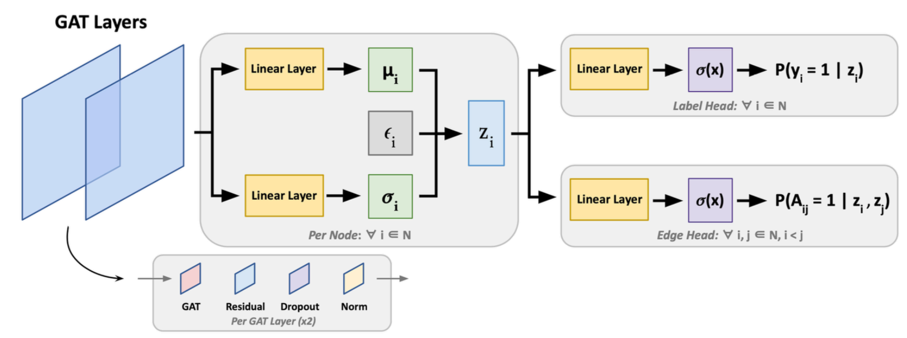

Figure 5: VGAE Architecture Diagram

4.2 Results

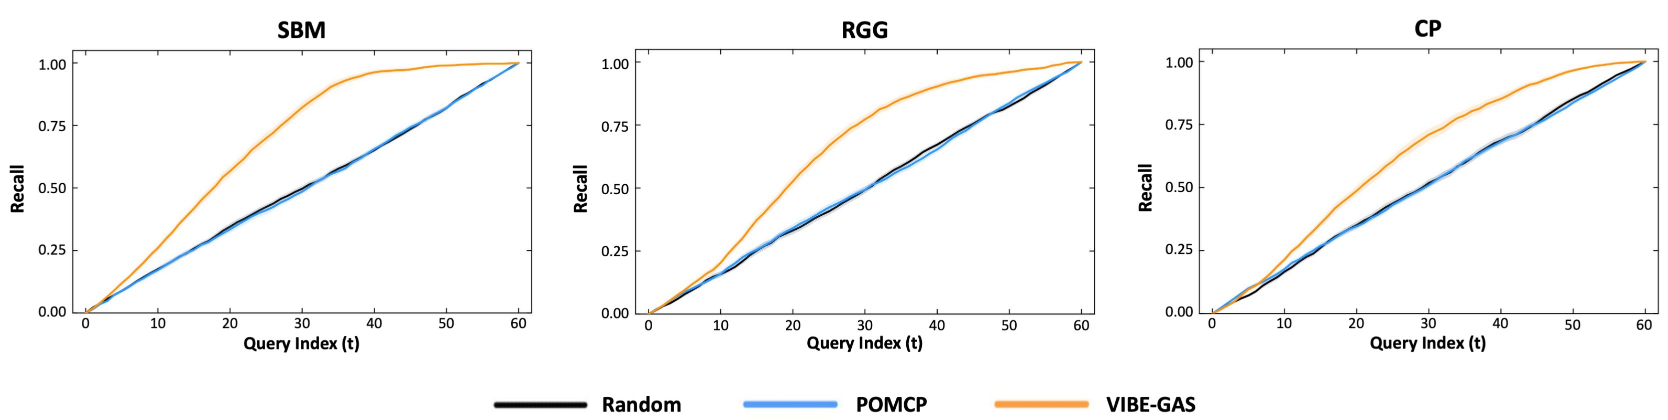

Figure 6: Positive Recalled Fraction Over an Episode

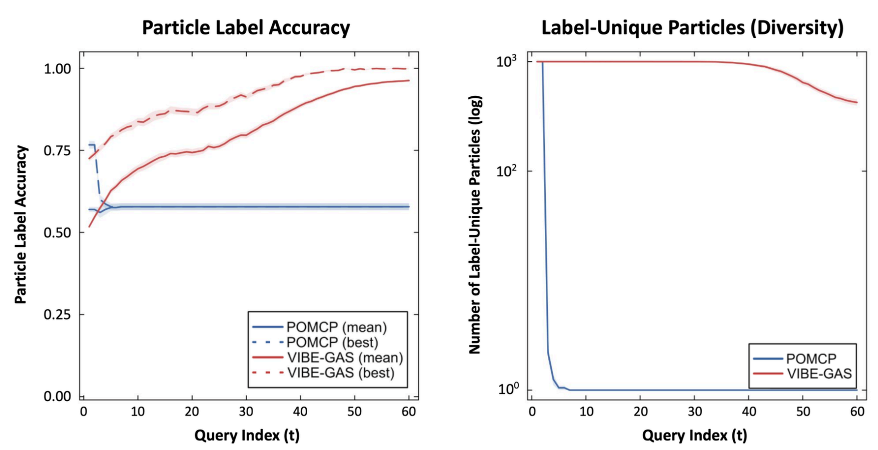

Figure 7: Belief Quality Analysis

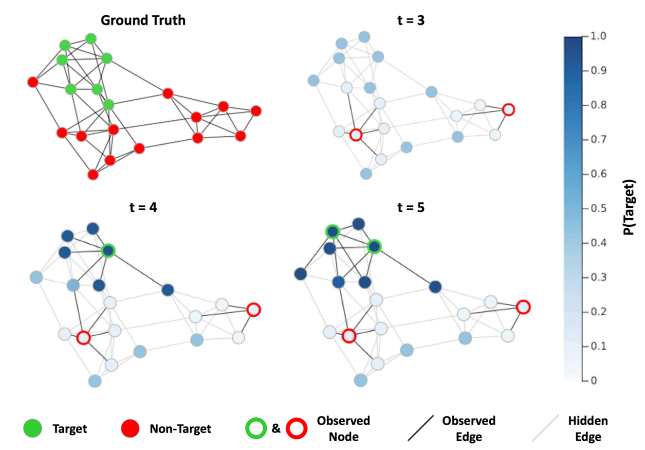

Figure 8: VGAE Belief Updates

5. Justifying VIBE-GAS

Training maximises the ELBO with respect to ϕ:

$$\mathcal{L} = \mathbb{E}_{Z\sim q_{\phi_{\text{enc}}}(Z|\tilde{A}_t,\tilde{Y}_t)}[\log p_{\phi_{\text{dec}}}(g \mid Z)] - \mathrm{KL}(q_{\phi_{\text{enc}}}(Z \mid \tilde{A}_t,\tilde{Y}_t) \parallel p(Z))$$

If the training data is representative of the test-time environment and posterior collapse is avoided, maximising the ELBO minimises an upper bound on the forward KL divergence between the true posterior p_t and variational estimate $b_{\phi,t}$ which can bound the 1-Wasserstein metric:

$$W_1(p_t, b_{\phi,t}) \leq D_{\max} \cdot \sqrt{\frac{1}{2}\mathrm{KL}(p_t \parallel b_{\phi,t})}$$

From here, the cumulative regret across an episode can be bounded as:

$$\mathrm{Regret}(\pi) \leq 2 \cdot \mathbb{E}_{h\sim\pi}\left[\sum_{t=0}^{H}\gamma^t \cdot C_{V_t} \cdot W_1(p_t, \hat{b}^K_{\phi,t})\right]$$

Hence, sensibly training a VGAE, directly reduces cumulative regret; effectively improving the induced policy and active search performance.

6. Conclusions

Primary Takeaways:

- VIBE-GAS leverages a VGAE in order to produce a conditioned particle belief that can be used for effective downstream planning.
- Standard particle filter bottlenecks such as particle degeneracy are structurally avoided; instead shifting sub-optimality to a controllable bias that is directly reduced through well-calibrated training.
- Empirically, VIBE-GAS outperforms both a random baseline as well as a standard particle filter POMCP planner in an active search setting.

Limitations:

- Computational complexity: VIBE-GAS scales quadratically with nodes (N) and linearly with particles (K). Formally; $O(KN^2)$ per step.
- Dependence on informative graph structure as well as the homophily assumption for the graph encoder to work effectively.

Future Work:

- Further empirical studies on real world data sets (e.g. CiteSeer).
- Improved, less coarse, and more generalised theoretical bounds.
- Full Deep RL approach to directly learn a policy over actions.

7. (Core) References

[1] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying, 2012.

[2] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains, 1998.

[3] David Silver and Joel Veness. Monte-Carlo planning in large pomdps, 2010.

[4] Arnaud Doucet and Adam Johansen. A tutorial on particle filtering and smoothing: Fifteen years later, 2009.

[5] Thomas N. Kipf and Max Welling. Variational graph auto-encoders, 2016.

Teim Giesen | *teimgiesen@tudelft.nl*

Research Project: CSE3000