

INFORMATION HOMEOSTASIS AS A NASH EQUILIBRIUM

A Potential Game for Decentralized, Multi-Agent Subterranean Exploration

Avi Soval, Keenan Albee



MOTIVATION

- Multi-agent, autonomous exploration is vital for search and rescue, distributed asset coverage, and planetary science
- Subterranean missions pose significant challenges in communication range, persistence, bandwidth, AoI
- As information parity is lost, a global belief state rapidly breaks down




FIG II (Top). The JPL CADRE mission will field a team of lunar rovers with a robust leader-election policy. [1]



FIG III (Left). Winning trial from team Cerberus at the DARPA SUBT challenge [2]

This work develops an approach to explore if *information homeostasis* - a steady global belief state - can be induced as the Nash equilibria of a potential game, such that consensus and convergence guarantees can be provided in the absence of any central coordinator

BACKGROUND

I. Multi-Agent Potential Games

- Potential games converge to Nash equilibria under pure best response dynamics [5,6], where any unilateral move is equivalent to a shift in the shared potential

$$U_i(a'_i; \mathbf{a}_{-i}) - U_i(a_i; \mathbf{a}_{-i}) = \Phi(a'_i; \mathbf{a}_{-i}) - \Phi(a_i; \mathbf{a}_{-i})$$

- Well-extended to multi-agent systems [7, 8, 9, 11]

II. Comms-Limited Belief Space Planning

- Several systems demonstrate successful coordination despite communication constraints - determining *when* and *how much* information to share is critical [12, 13, 14]
- Belief-space [15, 16, 17] planning using a representative, often computationally simpler world model is very well-suited to bandwidth- and info-limited missions

III. Bioelectric Signaling

- Scalar voltage signals propagating across gap junctions allow cells to broadcast their error from some setpoint
- High-level target states can be realized as dynamical attractors, and sharing of these 'stress' scalars is sufficient for decentralized goal alignment [3, 4]

KEY INVESTIGATIONS AND CONTRIBUTIONS

- Can multi-agent *exploration* be cast as a Potential Game?
- Can a state of global information parity be achieved purely by following best-response dynamics? Does a Nash equilibria exist?
- Can we quantify the 'potential-ness' [10] of our constructed game?
- Does an affine belief-coupling term break convergence guarantees?

$$\Phi(\mathbf{a}) = \sum_{f \in \mathcal{F}} \sum_{k=1}^{n_f(\mathbf{a})} (\alpha \text{IG}_{\pi_f(k)}(f) - \beta k) + \gamma \Psi(\mathbf{B})$$

$$\Psi(\mathbf{B}) = -\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} \|B_i - B_j\|_{1,\Omega_{ij}}$$

FIG I(d). Our formulation admits a canonical Potential game when driven by only greedy exploration and congestion. The goal is to quantify the impact of the affine belief-sharing term and its impact on convergence

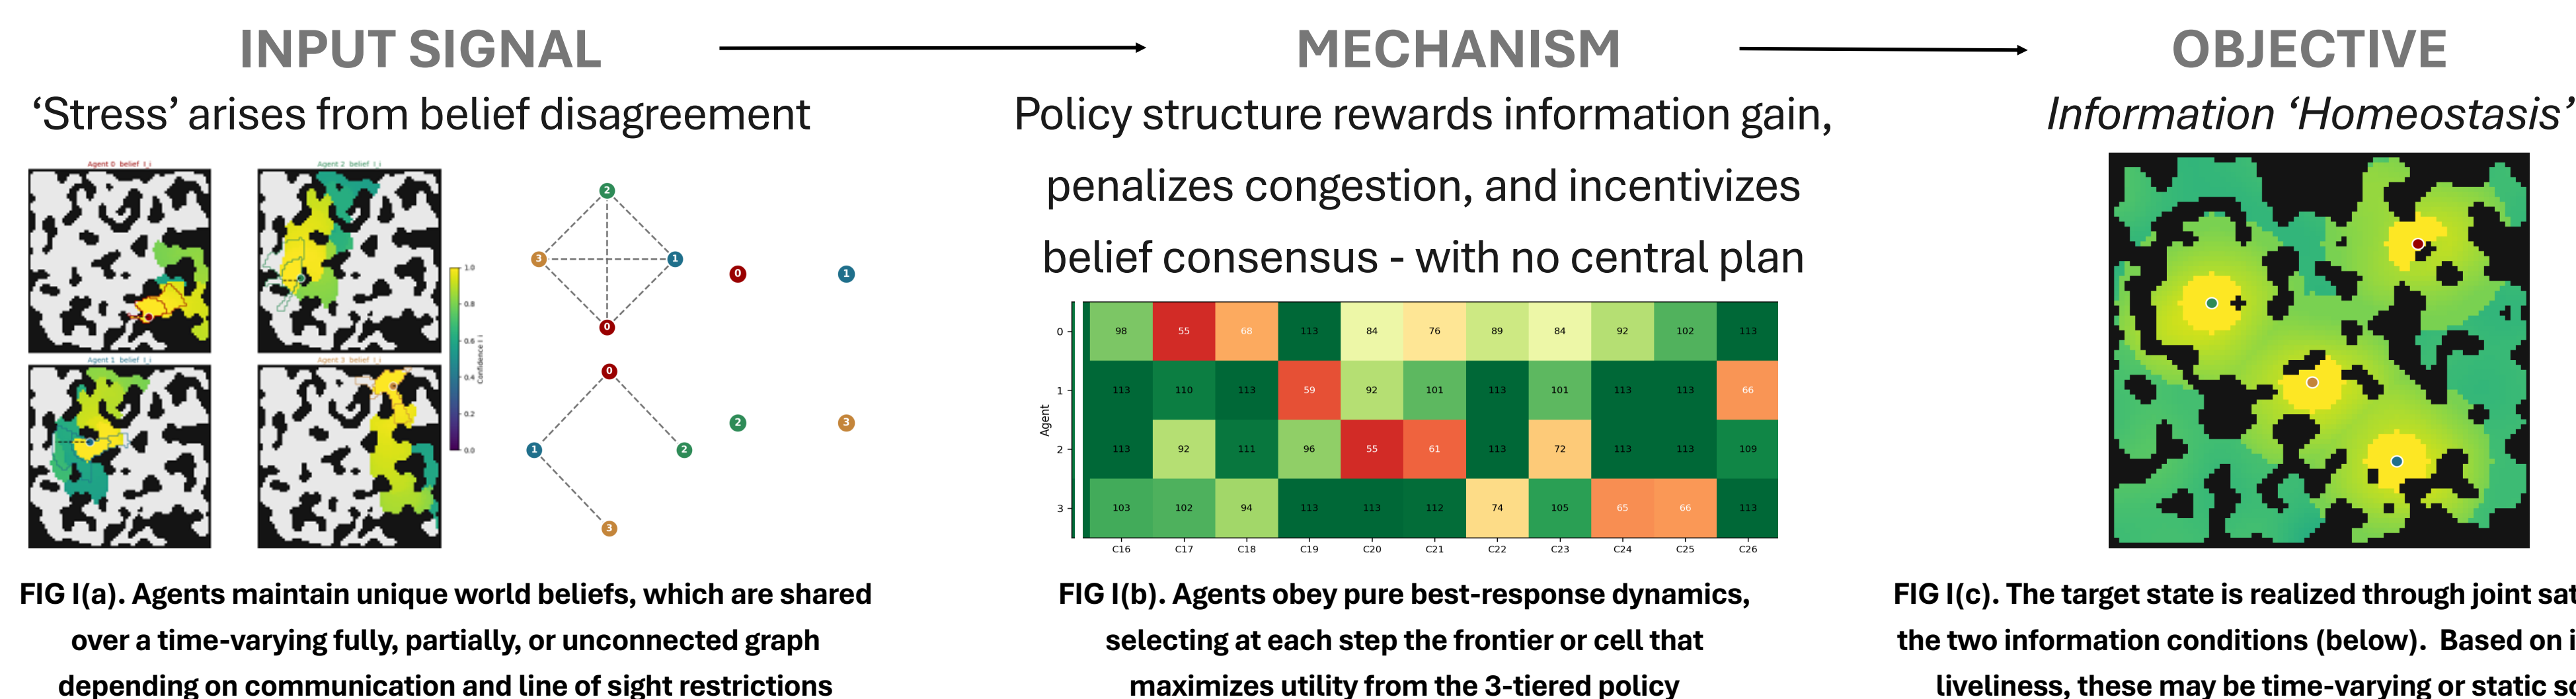


FIG I(a). Agents maintain unique world beliefs, which are shared over a time-varying fully, partially, or unconnected graph depending on communication and line of sight restrictions

FIG I(b). Agents obey pure best-response dynamics, selecting at each step the frontier or cell that maximizes utility from the 3-tiered policy

FIG I(c). The target state is realized through joint satisfaction of the two information conditions (below). Based on information liveliness, these may be time-varying or static solutions.

FORMULATION + METHODOLOGY

The Exploration Game




FIG IV. Initial setup with 4 agents, high-fidelity map data within sensing radius, comms over larger LOS

ENVIRONMENT | $G \in \{0,1\}^{H \times W}$

4-connected binary occupancy grid

PLAYERS | $\mathcal{N} = \{1, \dots, n\}$

Collect data within a sensing radius, and share data within a neighbor set determined by line of sight

RESOURCES | $\mathcal{F}(t)$

Time-varying frontier clusters

Objective: Information Homeostasis

RESIDUAL UNCERTAINTY

$$\bar{B} = \frac{1}{n} \sum_{i=1}^n B_i \rightarrow \frac{1}{|G_{\text{open}}|} \sum_{(r,c) \in G_{\text{open}}} \bar{B}(r,c) \leq \epsilon_B$$

Limits per-agent uncertainty in each cell

BELIEF CONSENSUS

$$\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} \|B_i - B_j\|_{1,\Omega_{ij}} \leq \epsilon_\sigma$$

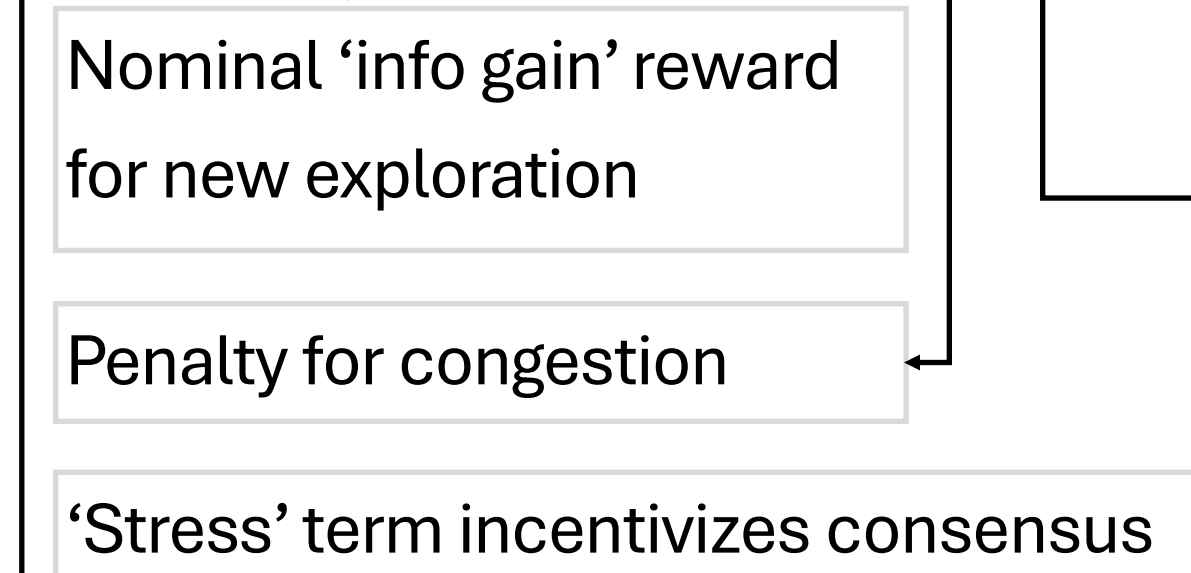
$$\|B_i - B_j\|_{1,\Omega_{ij}} := \frac{1}{|\Omega_{ij}|} \sum_{(r,c) \in \Omega_{ij}} |B_i(r,c) - B_j(r,c)|$$

Limits disagreement over all agent beliefs

Balancing Exploration + Consensus

UTILITY + ACTIONS | $a_i^* \in \arg \max U_i$

$$U_i(a_i; \mathbf{a}_{-i}, B_i) = \alpha \text{IG}_i(a_i) - \beta C_i(a_i; \mathbf{a}_{-i}) + \gamma \Delta \sigma_i(a_i)$$



Imperfect Information, and Confidence in Belief Sharing

Agents maintain a confidence in their belief, which can decay over time (for dynamic environments), and do not fully trust maps received from neighbors

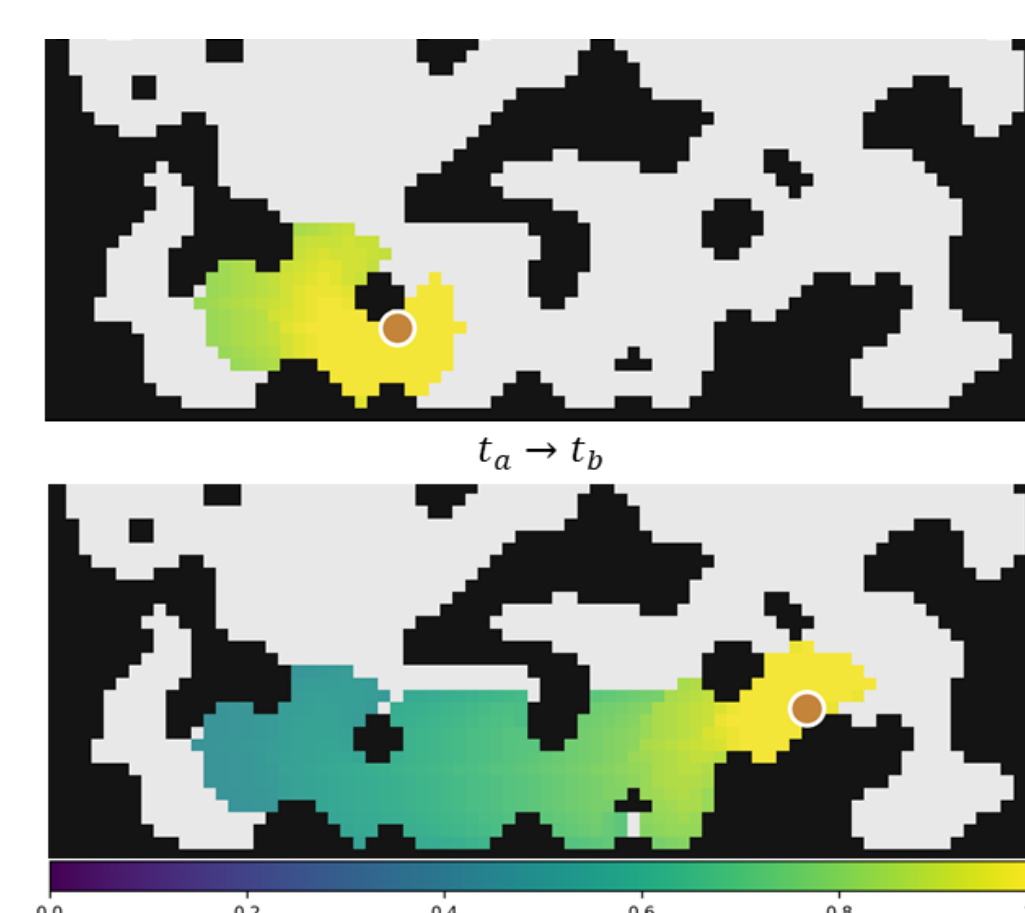



FIG V. Confidence in previously visited cells decays over time; cells in sensing range are 1.0

FIG VI. Agents #1 (red) and #2 (blue) have unique belief maps. Once LOS is established, maps are exchanged with a confidence penalty. Agent #1 does not believe everything it received about the map from #2.



INITIAL RESULTS

- Potential game implemented for the belief-less, ($\gamma=0$) uncoupled policy
- Used as a benchmark for potential-ness and for coverage time

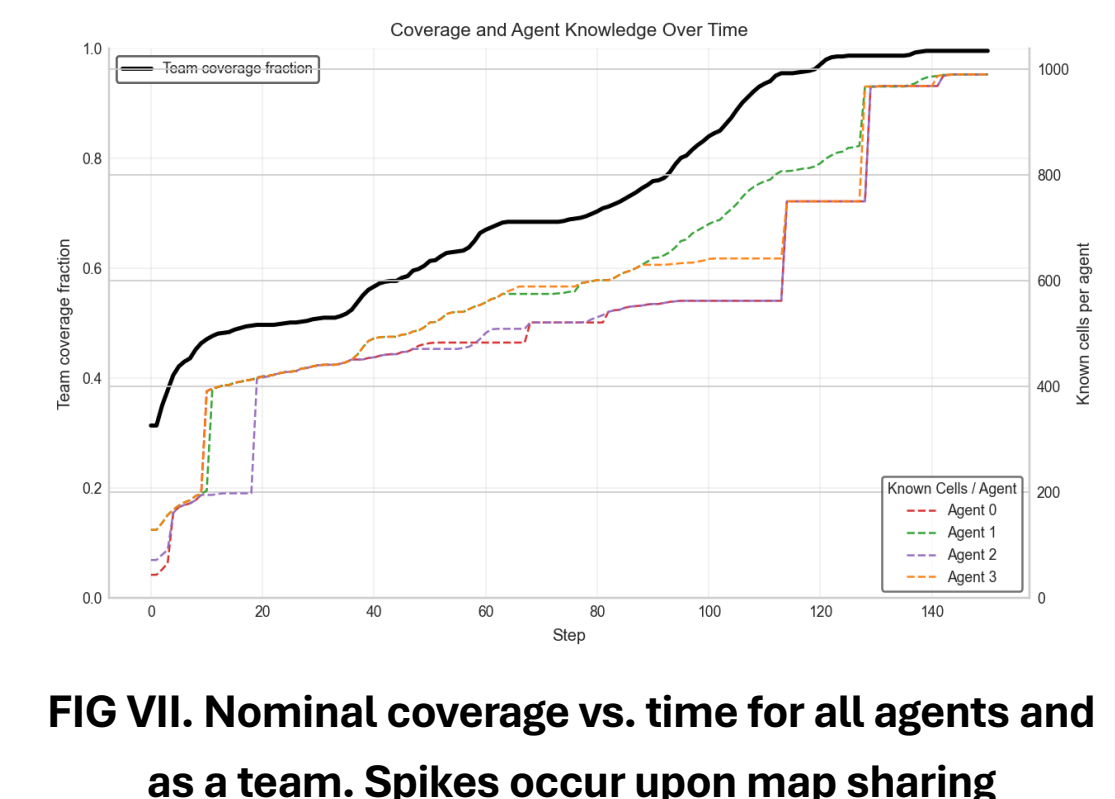


FIG VII. Nominal coverage vs. time for all agents and as a team. Spikes occur upon map sharing




FIG VIII. Weighting the belief-sharing too heavily leads agents to stay in place to maintain high confidence, rather than continue exploring

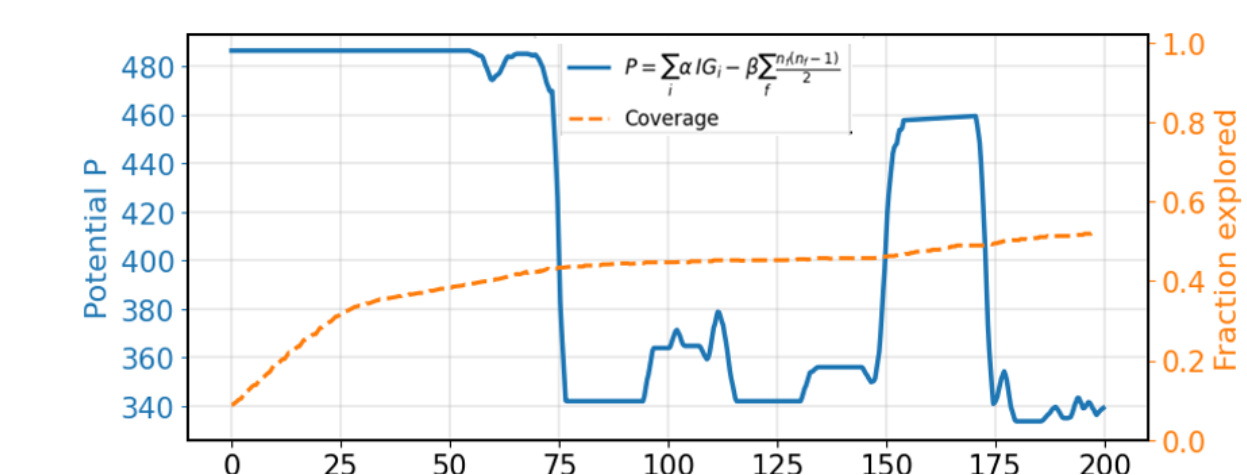


FIG IX. For the uncoupled game, the aggregate potential (from per-agent utilities) decreases over an interval

- Belief-consensus tuning underway; sensitive and can easily bias towards stationary coverage
- Preliminary results suggest the presence of a potential-like, always decreasing quantity (within a time-varying interval)

FUTURE WORK

- Fully implement the coupled, belief-sharing game
- Quantify 'potential-ness' of our perturbed game via combinatorial decomposition as in [10]
- Formally log residual potential, used to quantify any present convergence behavior or properties
- Monte-Carlo over weights, N agents, comms radii, map size
- Evaluate phases: dispersion, coordination, and homeostasis
- Restrict sharing bandwidth and enforce minimal signaling

REFERENCES

[1] S. Nayak, G. Lim, F. Rossi, M. Otte, and J.-P. de la Croix, "Multi-robot exploration for the CADRE mission," *Autonomous Robots*, vol. 49, no. 17, 2025.

[2] T. H. Chung, V. Orekhov, and A. Maio, "Into the robotic depths: Analysis and insights from the DARPA subterranean challenge," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, pp. 477-502, 2023.

[3] M. Levin and C. J. Martyniuk, "The bioelectric code: An ancient computational medium for dynamic control of growth and form," *BioSystems*, vol. 150, pp. 117-133, 2017.

[4] L. Shreesha and M. Levin, "Stress sharing as cognitive glue for collective intelligences: A computational model of stress as a coordinator for morphogenesis," *Biochemical and Biophysical Research Communications*, 2024.

[5] R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65-67, 1973.

[6] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124-143, 1996.

[7] J. R. Marden, G. Arslan, and J. S. Shamma, "Cooperative control and potential games," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 6, pp. 1393-1407, 2009.

[8] N. Hayashi, T. Ushio, and T. Kanazawa, "Potential game theoretic approach to power-aware mobile sensor coverage problem," *IEEE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, vol. E94-A, no. 3, pp. 929-936, 2011.

[9] Z. Williams, J. Chen, and N. Mehr, "Distributed potential IQR: Scalable game-theoretic trajectory planning for multi-agent interactions," 2023.

[10] M. Bichler, D. Legacci, P. Mertikopoulos, M. Oberbretner, and B. Pradel, "Characterizing the convergence of game dynamics via potentialness," 2025.

[11] A. Mandat et al., "Behaviorally heterogeneous multi-agent exploration using distributed task allocation," 2025.

[12] J. Gielis et al., "A critical review of communications in multi-robot systems," *Current Robotics Reports*, vol. 3, pp. 211-234, 2022.

[13] A.-a. Agha-mohammadi et al., "NeBula: TEAM CoSTAR's robotic autonomy solution that won phase II of DARPA subterranean challenge," *Field Robotics*, vol. 4, no. 1, pp. 601-649, 2024.

[14] M. V. Saboia da Silva et al., "ACHORD: Communication-aware multi-robot coordination with intermittent connectivity," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10184-10191, 2022.

[15] B. Franci and F. Fabiani, "Sharing beliefs to learn Nash equilibria," in 2024 European Control Conference (ECC), pp. 317-322, 2024.

[16] K. Pongsainginda, Z. Cao, B. P. L. Lau, R. Liu, C. Yuen, and U.-X. Tan, "Mel-explore: Communication-constrained multi-robot entropy-field-based exploration," *IEEE Transactions on Automation Science and Engineering*, vol. 22, p. 16062-16076, 2025.

[17] A. Patwardhan and A. J. Davison, "A distributed multi-robot framework for exploration, information acquisition and consensus," in 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 12062-12068, 2024.