

Bandwidth-constrained Variational Message Encoding for Cooperative Multi-agent Reinforcement Learning

Wei Duan

University of Technology Sydney
Sydney, Australia
wei.duan@uts.edu.au

En Yu

University of Technology Sydney
Sydney, Australia
en.yu-1@uts.edu.au

Jie Lu

University of Technology Sydney
Sydney, Australia
jie.lu@uts.edu.au

Junyu Xuan

University of Technology Sydney
Sydney, Australia
junyu.xuan@uts.edu.au

ABSTRACT

Graph-based multi-agent reinforcement learning (MARL) enables coordinated behavior under partial observability by modeling agents as nodes and communication links as edges. While recent methods excel at learning sparse coordination graphs—determining *who* communicates with *whom*—they do not address *what* information should be transmitted under hard bandwidth constraints. We study this bandwidth-limited regime and show that naïve dimensionality reduction consistently degrades coordination performance. Hard bandwidth constraints force selective encoding, but deterministic projections lack mechanisms to control how compression occurs. We introduce **Bandwidth-constrained Variational Message Encoding (BVME)**, a lightweight module that treats messages as samples from learned Gaussian posteriors regularized via KL divergence to an uninformative prior. BVME’s variational framework provides principled, tunable control over compression strength through interpretable hyperparameters, directly constraining the representations used for decision-making. Across SMACv1, SMACv2, and MPE benchmarks, BVME achieves comparable or superior performance while using 67–83% fewer message dimensions, with gains most pronounced on sparse graphs where message quality critically impacts coordination. Ablations reveal U-shaped sensitivity to bandwidth, with BVME excelling at extreme ratios while adding minimal overhead.

KEYWORDS

Multi-agent reinforcement learning, Graph-based coordination, Variational message encoding, Bandwidth-limited communication

ACM Reference Format:

Wei Duan, Jie Lu, En Yu, and Junyu Xuan. 2026. Bandwidth-constrained Variational Message Encoding for Cooperative Multi-agent Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 10 pages. <https://doi.org/10.65109/QXVZ8292>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/QXVZ8292>

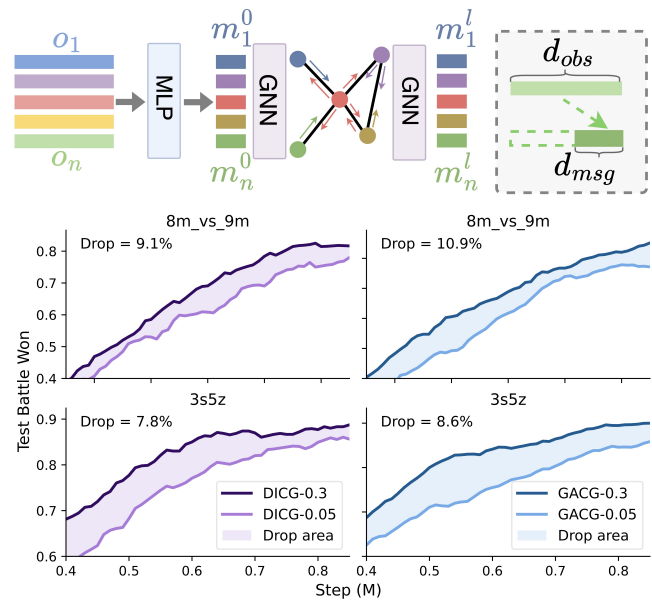


Figure 1: Bandwidth constraints degrade performance. Top: Standard GNN-based MARL encodes observations $o_i \in \mathbb{R}^{d_{obs}}$ into messages $m_i^l \in \mathbb{R}^{d_{msg}}$. Compression ratio $r = d_{msg}/d_{obs}$ controls bandwidth. Bottom: Test win rates on SMACv1 when reducing r from 0.3 to 0.05. Shaded regions show performance loss (drop area).

1 INTRODUCTION

Cooperative multi-agent reinforcement learning (MARL) under partial observability fundamentally requires information sharing among agents [7, 12, 21, 34, 39]. Graph-based formulations, where agents are nodes and edges represent coordination dependencies, have become a dominant framework for scalable cooperation [2, 15, 20, 37, 40, 42, 43]. When combined with graph neural networks (GNNs) [5, 9, 10, 19], these approaches enable end-to-end learning of coordination policies through differentiable message passing, achieving strong performance in domains ranging from cooperative navigation to strategy games [13, 35, 37]. Despite this success, a fundamental challenge remains underexplored: how to optimally encode information when communication capacity is limited.

Existing work addresses two complementary aspects of multi-agent communication. **Topology-focused methods** learn **who** should communicate: attention mechanisms [14, 16, 17, 20, 22, 29] dynamically weight interactions, while sparse coordination graphs [6, 8] prune unnecessary connections. These approaches implicitly assume that once a link is established, sufficient bandwidth exists to transmit high-dimensional features. In contrast, **content-focused methods** [3, 12, 30, 31, 36] address **when** and **what** to communicate through gating and information-theoretic regularization, but typically omit explicit graph structure. Recent work has explored multi-level communication [44] and interpretable information gating [1], while MAGI [4] integrates information bottlenecks into GNNs layers to filter noisy channels. However, these approaches share a common limitation: they do not explicitly model hard bandwidth limits—scenarios where message dimensionality is strictly bounded (e.g., 5–10% of observation size) due to physical or scalability constraints. MAGI [4], for instance, relies on dense attention mechanisms rather than learned sparse graphs, while information-theoretic methods typically regulate soft capacity rather than enforcing dimensional constraints.

As illustrated in Fig. 1 (top), standard GNN-based MARL encodes each agent’s observation $o_i \in \mathbb{R}^{d_{\text{obs}}}$ through an MLP into an initial message m_i^0 , which is refined via message-passing layers into $m_i^l \in \mathbb{R}^{d_{\text{msg}}}$. The compression ratio $r = d_{\text{msg}}/d_{\text{obs}}$ defines communication bandwidth. Existing methods [6, 8, 20] apply deterministic linear projections to satisfy dimensional constraints, but when r is small (≤ 0.1), this approach discards information indiscriminately. To quantify this effect, we conduct a controlled study on DICG [20] (dense) and GACG [6] (sparse) using SMACv1 [28], progressively reducing r from 0.3 to 0.05. Both suffer consistent performance drops, with sparse graphs exhibiting sharper degradation (Fig. 1, bottom). We quantify this via *drop area*—the cumulative loss between bandwidth curves. Sparse graphs are more sensitive because each edge must carry more task-critical information, demonstrating that topology and bandwidth interact in task-dependent ways.

Such bandwidth constraints arise naturally in real-world deployments, such as warehouse robot swarms that train centrally using full state observations but execute via local wireless radios with limited throughput. In these scenarios, hard bandwidth constraints force selective encoding, yet deterministic projections lack mechanisms to control how compression occurs—they discard information uniformly regardless of task relevance. This motivates our central question: **How can agents achieve effective coordination when communication bandwidth is severely limited?**

To address this question, we propose **Bandwidth-constrained Variational Message Encoding (BVME)**. BVME models each agent’s message as a sample from a learned Gaussian distribution, regularized toward an uninformative prior via Kullback-Leibler (KL) divergence. This variational framework provides principled, tunable control over compression through interpretable hyperparameters, enabling: (1) direct regulation of how much information agents transmit via the KL penalty; (2) end-to-end learning of task-relevant representations under bandwidth constraints; (3) flexible adjustment of compression strength without architecture changes. A critical design element is *on-path* coupling: BVME directly injects sampled messages into each agent’s Q -network, ensuring the

KL regularization operates on the same representations that drive coordination decisions, rather than filtering auxiliary signals that bypass decision-making.

We instantiate BVME on GACG [6], which learns sparse relations between agents and groups, demonstrating how variational encoding complements structured topologies. We choose SMACv1 [28], SMACv2, and MPE-Tag to evaluate our method, comparing against QMIX [27], DICG [20], and GACG [6]. Under severe bandwidth constraints (5%), BVME achieves higher win rates and faster convergence. Ablations reveal a U-shaped performance curve: substantial gains at low bandwidth, marginal benefits at moderate compression, and neutral impact at high capacity. On-path coupling proves essential, substantially outperforming off-path regularization.

Our contributions are:

- (1) We propose BVME, a variational framework providing **principled, tunable control over message compression** through KL regularization. Unlike deterministic projections that discard information uniformly, BVME enables explicit capacity control via interpretable hyperparameters (r , σ_0 , λ_{KL}).
- (2) We identify a U-shaped sensitivity to bandwidth: BVME excels at extreme compression ($r \leq 0.05$) where it filters noise and prioritizes coordination-critical features, with on-path coupling essential for effectiveness.
- (3) Experiments across SMACv1/v2 and MPE-Tag show BVME enables stable learning where baselines fail (5% bandwidth) and achieves comparable performance with 67–83% fewer message dimensions on sparse graphs.

2 RELATED WORK

2.1 Graph-structured coordination in MARL

Graph neural networks have become a dominant framework for cooperative MARL under partial observability [2, 15], often combined with value decomposition [26, 27, 32] for decentralized execution with centralized training. Methods for learning *who* communicates can be categorized by their induced communication topology.

Dense communication. DCG [2] and DICG [20] use complete coordination graphs or dense neural networks to model all agent interactions, while attention-based approaches like ATOC [16] and MAAC [14] apply learned weights across the full team. Though adaptively weighted, these methods scale poorly as team size grows. **Sparse coordination graphs.** SOP-CG [38], GACG [6], LTS-CG [8], and CommFormer [13] learn sparse topologies through value decomposition over restricted structures, group-aware distributions, temporal inference, or bi-level optimization, enabling scalable coordination by restricting communication to task-relevant subsets.

Both paradigms optimize *who* communicates but treat messages as deterministic embeddings, overlooking *what* information should be transmitted under bandwidth constraints. When message dimensionality is limited, even well-structured topologies may propagate redundant or noisy content.

2.2 Communication-efficient MARL

Complementary research explores *when* and *what* to communicate, often without explicit graph structure. IC3Net [30] and TarMAC [3] learn dynamic communication policies, while information-theoretic

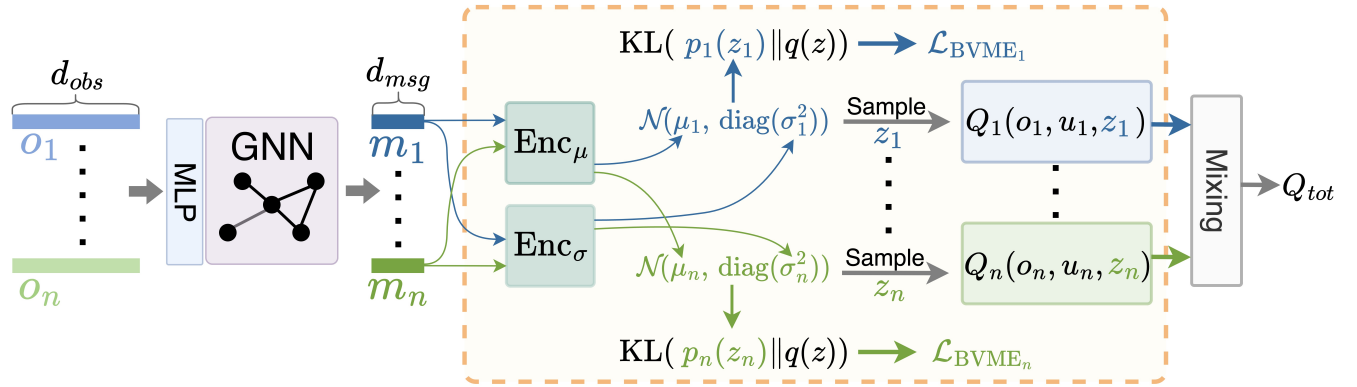


Figure 2: BVME architecture. *Standard approach (left):* GNN-based MARL encodes observations $\{o_i\}$ via MLP compression followed by graph convolution into messages $\{m_i\}$ that directly feed Q -functions. *BVME (right, orange box):* Each m_i parameterizes a Gaussian posterior via encoders Enc_μ and Enc_σ . Sampled messages z_i are used for Q_i estimation, while KL divergence to a prior $q(z) = \mathcal{N}(0, \sigma_0^2 I)$ enforces bandwidth constraints. The compression ratio $r = d_{\text{msg}}/d_{\text{obs}}$ controls message dimensionality. Training combines TD loss with BVME regularization.

methods like NDQ [36] promote task-relevant messaging by penalizing unnecessary entropy. Recent work explores hierarchical [44] and interpretable [1] communication mechanisms. Liu & Bai [21] provide theoretical analysis of information sharing benefits under partial observability. However, these methods typically lack structural biases that enhance scalability and credit assignment.

Prior work on bandwidth-limited MARL addresses orthogonal aspects of communication efficiency. Scheduling methods [17, 24] learn *when/who* communicates via gating or Top- k selection, reducing *frequency* without compressing *content*. MAGI [4] applies information bottlenecks to GNN stacks but targets adversarial robustness via dense attention rather than dimensional constraints with sparse graphs. Existing methods treat bandwidth heuristically through projections or auxiliary losses—without explicit control over the message dimensionality that directly informs decisions. BVME addresses this via on-path KL regularization, ensuring bandwidth limits constrain representations that drive coordination.

3 PRELIMINARIES

3.1 Cooperative Dec-POMDPs

Cooperative multi-agent tasks can be modeled as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) [25] with tuple $\langle \mathcal{A}, \mathcal{S}, \{\mathcal{U}_i\}_{i=1}^n, P, \{O_i\}_{i=1}^n, \{\pi_i\}_{i=1}^n, R, \gamma \rangle$. At time t , each agent $i \in \mathcal{A} = \{1, \dots, n\}$ receives a local observation $o_i^t \in O_i$, maintains an action-observation history $\tau_i^t = (o_i^0, u_i^0, \dots, o_i^t)$, and selects an action $u_i^t \in \mathcal{U}_i$ via a stochastic policy $\pi_i(u_i^t | \tau_i^t)$. The joint action $\mathbf{u}^t = (u_1^t, \dots, u_n^t)$ induces the next state $s^{t+1} \sim P(\cdot | s^t, \mathbf{u}^t)$ and yields a shared reward $r^t = R(s^t, \mathbf{u}^t)$. The objective is to maximize the expected discounted return:

$$Q_{\text{tot}}(s, \mathbf{u}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r^t \mid s^0 = s, \mathbf{u}^0 = \mathbf{u} \right], \quad \gamma \in [0, 1). \quad (1)$$

We adopt the CTDE paradigm [27, 41]: during training, a central learner accesses global state s^t and all observations $\{o_i^t\}_{i=1}^n$ to optimize policies and coordination structures; during execution, each

agent independently selects actions via $\pi_i(\tau_i^t)$ using only local history and received messages. Bandwidth constraints $r = d_{\text{msg}}/d_{\text{obs}}$ on inter-agent communication ensure deployability on resource-constrained platforms. A widely used approach is *value decomposition*, where per-agent $Q_i(\tau_i, u_i)$ are combined to calculate Q_{tot} .

3.2 Graph Neural Networks for Agent Communication

Graph neural networks (GNNs) are widely used in cooperative MARL to support inter-agent communication via a *coordination graph* whose nodes are agents and edges represent communication links. Given an adjacency matrix $\tilde{A} \in \mathbb{R}^{n \times n}$ and agent features $X \in \mathbb{R}^{n \times d}$, a message-passing layer computes

$$M = \phi(\tilde{A}XW), \quad (2)$$

where W is a learnable weight matrix and $\phi(\cdot)$ is a nonlinearity. After L layers, the final messages $M^{(L)}$ are integrated into per-agent:

$$Q_i(\tau_i, u_i) = f_{\theta_i}(o_i^t, u_i, m_i^{(L)}), \quad (3)$$

where $m_i^{(L)}$ is the i -th row of $M^{(L)}$. The utilities $\{Q_i\}_{i=1}^n$ are then combined via value decomposition to produce Q_{tot} .

3.3 Group-Aware Coordination Graphs (GACG)

A key challenge in GNN-based MARL is determining the adjacency \tilde{A} that governs message passing. We build on GACG [6], which learns sparse, group-aware coordination graphs by representing the adjacency as a Gaussian distribution:

$$\text{vec}(\tilde{A}^t) \sim \mathcal{N}(\mu_A^t, \Sigma_A^t), \quad (4)$$

where μ_A^t encodes pairwise agent importance derived from current observations, and Σ_A^t captures group-level correlations estimated from recent trajectory history. This formulation enables both sparse edge selection and discovery of coordinated sub-teams.

GACG is trained with a QMIX-style TD loss and a group regularizer that encourages within-group cohesion and between-group

specialization:

$$\mathcal{L}_{\text{GACG}} = \mathcal{L}_{\text{TD}} + \lambda_g \mathcal{L}_g, \quad (5)$$

where \mathcal{L}_g compares intra-group policy similarity to inter-group diversity (see [6] for details). In our work, we adopt GACG’s graph learning mechanism to obtain \tilde{A}^t , and focus on optimizing the message content M that flows on this learned sparse topology.

While GACG optimizes *which* edges to activate, it treats the messages M transmitted on these edges as deterministic embeddings. In the next section, we introduce BVME, which addresses *what* information should be encoded into M when message dimensionality is constrained.

4 METHOD: BANDWIDTH-CONSTRAINED VARIATIONAL MESSAGE ENCODING

Graph-structured coordination in MARL typically exchanges full or lightly projected messages between agents. When bandwidth is tight ($d_{\text{msg}} \ll d_{\text{obs}}$), naïve linear projection degrades coordination performance (Fig. 1). We propose *Bandwidth-constrained Variational Message Encoding* (BVME), which treats messages as samples from learned Gaussian posteriors regularized via KL divergence [18]. BVME provides principled, tunable control over compression strength through: compression ratio $r = d_{\text{msg}}/d_{\text{obs}}$, prior scale σ_0 , and KL weight λ_{KL} , enabling agents to learn compact representations under strict bandwidth limits.

4.1 Standard Message Passing

Given a learned adjacency \tilde{A} (Sec. 3.3) and agent features $X^{(l-1)} \in \mathbb{R}^{n \times d_{l-1}}$, a standard GNN message-passing layer computes:

$$M^{(l)} = \phi_l(\tilde{A}X^{(l-1)}W_l), \quad (6)$$

where $W_l \in \mathbb{R}^{d_{l-1} \times d_l}$ is a learnable weight matrix and ϕ_l is a non-linearity. To satisfy a bandwidth budget $r = d_{\text{msg}}/d_{\text{obs}}$, one sets the final layer width $d_l = d_{\text{msg}}$. This projection respects the dimensional constraint but is task-agnostic: when $r \leq 0.3$, performance drops sharply (Fig. 1).

4.2 Variational Message Encoding

From Hard Constraints to Tunable Compression Control. Hard bandwidth constraints force selective encoding, but deterministic projections lack mechanisms to control how compression occurs—they discard information uniformly regardless of task relevance. By sampling messages from learned distributions regularized via KL divergence to an uninformative prior (Sec. 4.3), the KL penalty directly controls how much information each agent transmits. **Crucially, agents can adapt compression strength through tunable hyperparameters ($\sigma_0, \lambda_{\text{KL}}$), unlike fixed projections that compress uniformly regardless of coordination needs.** This enables: (1) explicit capacity control via tunable hyperparameters ($\sigma_0, \lambda_{\text{KL}}$); (2) end-to-end learning of task-relevant representations under bandwidth constraints; (3) stochastic regularization that may improve generalization.

Gaussian parameterization. Rather than directly using the GNN outputs $M^{(l)} = \{m_1^{(l)}, \dots, m_n^{(l)}\}$ from Eq. (6), we use each $m_i^{(l)} \in \mathbb{R}^{d_{\text{msg}}}$ to parameterize a diagonal Gaussian posterior via

lightweight encoders:

$$\mu_i^{(l)} = \text{Enc}_\mu(m_i^{(l)}), \quad \log \sigma_i^{2(l)} = \text{Enc}_\sigma(m_i^{(l)}), \quad (7)$$

where $\text{Enc}_\mu, \text{Enc}_\sigma : \mathbb{R}^{d_{\text{msg}}} \rightarrow \mathbb{R}^{d_{\text{msg}}}$ are single-layer MLPs. These encoders map each GNN message to the mean and log-variance of a distribution, enabling the model to learn which message features are certain (low variance) versus uncertain (high variance). Agent i ’s message distribution is then:

$$p_i^{(l)}(z) = \mathcal{N}(z \mid \mu_i^{(l)}, \text{diag}(\sigma_i^{2(l)})). \quad (8)$$

Each agent samples its Bandwidth-constrained message via reparameterization:

$$z_i^{(l)} = \mu_i^{(l)} + \sigma_i^{(l)} \odot \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, I), \quad (9)$$

where \odot denotes element-wise multiplication. This standard reparameterization trick allows gradients to flow through the stochastic sampling operation. Stacking sampled messages across agents yields:

$$Z^{(l)} = [z_1^{(l)}, \dots, z_n^{(l)}]^\top \in \mathbb{R}^{n \times d_{\text{msg}}}, \quad (10)$$

which replaces $M^{(l)}$ in GNN layers or as input to Q -functions.

On-path sampling and stability. The reparameterization in Eq. (9) enables gradients to flow through $\mu_i^{(l)}$ and $\sigma_i^{(l)}$. To enforce bandwidth constraints, we will regularize these posteriors via KL divergence to an uninformative prior (Sec. 4.3). Crucially, we propagate the *sampled* $z_i^{(l)}$ —not its mean $\mu_i^{(l)}$ —into the Q -functions for value estimation and policy updates. This *on-path* coupling ensures the bandwidth regularization directly constrains the stochastic representations that drive decision-making, rather than regularizing an auxiliary pathway that bypasses the Q -network. In contrast, *off-path* approaches apply bottlenecks to separate encoder branches whose outputs do not directly inform control, weakening the connection between compression and coordination performance (see Sec. 5.3 for empirical validation).

To prevent variance collapse (where $\sigma_i^{(l)} \rightarrow 0$ eliminates stochasticity and degrades exploration), we clamp the log-variance:

$$\log \sigma_i^{2(l)} \in [\log \sigma_{\min}^2, \log \sigma_{\max}^2], \quad (11)$$

with $\sigma_{\min}=0.01$ and $\sigma_{\max} \in [1, 3]$. This maintains a minimum level of noise in the messages, encouraging the model to preserve uncertainty rather than collapsing to deterministic point estimates.

4.3 Bandwidth Enforcement via KL Divergence

Stochastic sampling alone does not guarantee compression—the model could set $\sigma_i^{(l)} \rightarrow 0$ and encode arbitrarily detailed features in $\mu_i^{(l)}$. To enforce meaningful compression, we regularize each agent’s message distribution following the information bottleneck principle [33]: limit transmitted information while preserving task-relevant content.

We achieve this by regularizing each agent’s posterior $p_i^{(l)}$ toward a shared uninformative prior via KL divergence. The prior serves as a reference: closer distributions transmit less information, defined as:

$$q(z) = \mathcal{N}(0, \sigma_0^2 I), \quad (12)$$

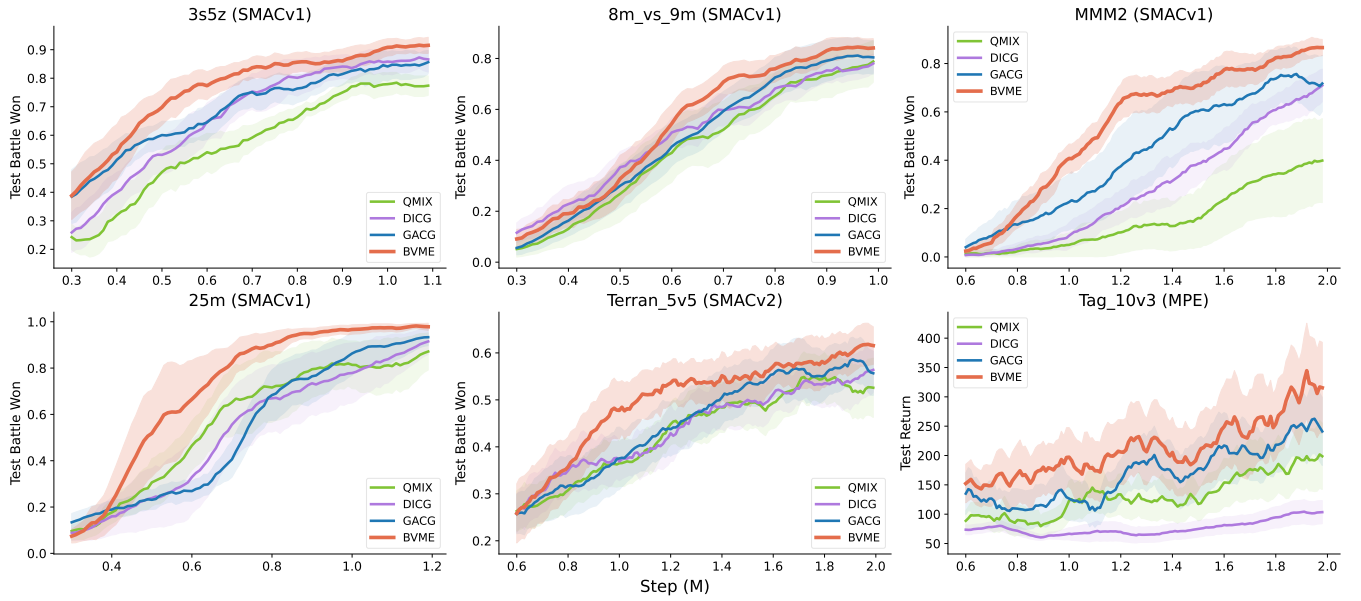


Figure 3: Overall performance across benchmarks. Learning curves on SMACv1 (3s5z, 8m_vs_9m, MMM2, 25m), SMACv2 (Terran_5v5), and MPE Tag at bandwidth ratio $r = 0.05$. BVME consistently outperforms QMIX, DICG, and GACG across all tasks, with shaded regions indicating standard error over 5 seeds.

a zero-mean Gaussian carrying no task-relevant information—agents transmitting messages sampled from $q(z)$ would communicate nothing useful. The KL divergence $\text{KL}(p_i^{(l)} \parallel q)$ measures how different agent i 's message distribution is from this uninformative prior. By penalizing large KL values, we limit how much agents can deviate from $q(z)$, directly controlling communication capacity. Note that $q(z)$ has no agent subscript—all agents share the same prior, ensuring identical bandwidth constraints. When σ_0 is very small, $q(z)$ approaches a point mass at zero, representing the limit of "transmitting nothing."

Why this enforces bandwidth limits. The prior variance σ_0^2 directly controls the information capacity of each message dimension. To see this concretely, consider how the KL penalty responds to different encoding strategies. If $\sigma_0 = 0.1$ (tight capacity), encoding a feature with $\mu_{i,d} = 5$ and $\sigma_{i,d} = 1$ yields a KL contribution of approximately $\frac{25+1}{0.01} - 1 \approx 2599$, creating strong pressure to reduce $\mu_{i,d}$ or eliminate this dimension. Conversely, with $\sigma_0 = 1$ (relaxed capacity), the same configuration yields $\frac{25+1}{1} - 1 = 25$, allowing more information transmission.

For a diagonal Gaussian posterior $p_i^{(l)} = \mathcal{N}(\mu_i^{(l)}, \text{diag}(\sigma_{i,d}^{2(l)}))$, the KL divergence has a closed form:

$$\text{KL}(p_i^{(l)} \parallel q) = \frac{1}{2} \sum_{d=1}^{d_{\text{msg}}} \left(\frac{\sigma_{i,d}^2 + \mu_{i,d}^2}{\sigma_0^2} - 1 + \log \frac{\sigma_0^2}{\sigma_{i,d}^2} \right), \quad (13)$$

where $\mu_{i,d}$ and $\sigma_{i,d}$ index dimension d of agent i 's posterior. This penalizes both large means ($\mu_{i,d}^2$ term) and variances ($\sigma_{i,d}^2$ term) relative to the prior, encouraging compact representations.

4.4 Training Objective

At each timestep t , each agent i samples a Bandwidth-constrained message z_i^t via Eq. (9) applied at the final GNN layer L . The local utility then conditions on this sampled message $Q_i(\tau_i^t, u_i^t, z_i^t)$, which captures how agent i evaluates action u_i^t given its history τ_i^t and the compressed message z_i^t . The mixer aggregates $\{Q_i\}_{i=1}^n$ into Q_{tot} following standard value decomposition (Sec. 3.3).

Let \mathcal{B} denote a replay batch with timestep mask $m_t \in \{0, 1\}$ indicating valid transitions. The BVME regularization term averages KL divergences across agents and timesteps:

$$\mathcal{L}_{\text{BVME}} = \frac{\lambda_{\text{KL}}}{n T_{\mathcal{B}}} \sum_{t \in \mathcal{B}} m_t \sum_{i=1}^n \text{KL}(p_i^{(L,t)} \parallel q), \quad (14)$$

where $T_{\mathcal{B}} = \sum_{t \in \mathcal{B}} m_t$ counts valid timesteps and $p_i^{(L,t)}$ denotes agent i 's posterior at layer L and time t . The KL weight λ_{KL} scales the regularization pressure, trading off compression against coordination performance. Together with the prior scale σ_0 (Eq. 12), these hyperparameters provide complementary control over bandwidth: σ_0 sets the baseline capacity per dimension, while λ_{KL} determines how strongly to enforce this constraint. While BVME can be applied at any GNN layer, we apply it only at the final layer L whose outputs feed the Q -functions, as this directly regularizes the representations used for control.

The full training objective combines the BVME regularization with the GACG backbone:

$$\mathcal{L} = \underbrace{\mathcal{L}_{\text{TD}} + \lambda_g \mathcal{L}_g}_{\text{GACG objective}} + \mathcal{L}_{\text{BVME}}. \quad (15)$$

During training, we sample stochastic messages via Eq. (9) to enable exploration and ensure the KL penalty shapes the representations.




Figure 4: U-shaped performance curve. Relative cumulative performance $\Delta AUC(r) = AUC_{BVME} - AUC_{GACG}$ across compression ratios. Green regions indicate BVME advantage; orange regions indicate GACG advantage. BVME excels at low (≤ 0.05) and high (≥ 0.15) bandwidths, with GACG competitive only in the mid-range (0.10–0.15).

At evaluation, we use the deterministic mean $z_i = \mu_i^{(L)}$ to reduce variance and improve stability.

5 EXPERIMENTS

We design our experiments to answer the following key questions: (1) How does BVME perform compared to MARL baselines across diverse benchmarks? (2) How sensitive is performance to the bandwidth ratio $r = d_{msg}/d_{obs}$? (3) What is the effect of feeding *sampled* variational messages into the Q -network (on-path) versus using only the mean (off-path)? (4) Is BVME particularly beneficial when combined with sparse-graph backbones (e.g., GACG), where fewer edges make message quality more critical? (5) How do the regularization strength λ_{KL} and prior scale σ_0 jointly affect performance?

All experiments are conducted on three cooperative MARL benchmarks with results averaged over 5 random seeds. **SMACv1** [28] provides standardized StarCraft II micromanagement scenarios, serving as a widely used benchmark for value-decomposition and graph-based MARL. **SMACv2** [11] increases stochasticity by randomizing start positions and unit types, testing robustness and generalization. **Tag** [23] emphasizes communication under partial observability in a predator-prey scenario.

5.1 Overall Performance

We compare BVME against representative MARL baselines:

- **QMIX** [27]: value decomposition with centralized training but no graph-structured communication.
- **DICG** [20]: fully connected coordination graph with attention-based message weighting.
- **GACG** [6]: state-of-the-art sparse coordination graph learning with group-aware structure discovery.
- **BVME (ours)**: variational Bandwidth-constrained message encoding applied to the GACG backbone.

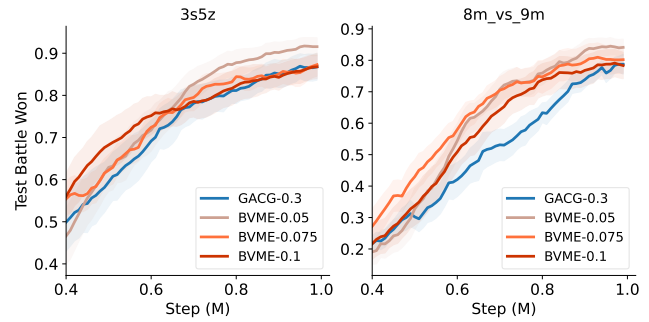


Figure 5: Extreme compression comparison. GACG at full budget ($r=0.30$) versus BVME at tighter ratios ($r=0.05, 0.075, 0.10$) on 3s5z and 8m_vs_9m. Even with 67–83% fewer message dimensions, BVME achieves comparable or superior win rates.

Map	Method	r	Msg Red. vs. 0.30	AUC	Win
3s5z	GACG-0.30	0.30	–	0.5201	0.8677
	BVME-0.05	0.05	83.3%	0.5517	0.9156
	BVME-0.075	0.075	75.0%	0.5427	0.8729
	BVME-0.10	0.10	66.7%	0.5570	0.8677
8m_vs_9m	GACG-0.30	0.30	–	0.3273	0.7875
	BVME-0.05	0.05	83.3%	0.3791	0.8411
	BVME-0.075	0.075	75.0%	0.4005	0.8021
	BVME-0.10	0.10	66.7%	0.3621	0.7833

Table 1: Comparison of GACG at full budget ($r = 0.30$) with BVME at reduced ratios on SMAC maps. “Msg Red. vs. 0.30” denotes the relative message dimensionality reduction compared to the $r = 0.30$ baseline, computed as $(0.30 - r)/0.30$.

Results. Fig. 3 shows learning curves across all benchmarks. BVME consistently outperforms baselines under bandwidth constraints ($r = 0.05$): on small-scale maps (3s5z, 8m_vs_9m), BVME achieves 5–10% higher final win rates than GACG, while on larger scenarios (25m, MMM2), it demonstrates both faster convergence and superior asymptotic performance. On SMACv2 (Terran_5v5), BVME maintains a clear advantage over both GACG and DICG despite increased stochasticity, demonstrating robustness to harder, randomized environments. On MPE Tag, BVME improves coordination in asymmetric predator-prey interactions, highlighting adaptability beyond StarCraft domains. These results establish BVME as an effective mechanism for multi-agent coordination under stringent bandwidth budgets, with benefits most pronounced on sparse graphs where message quality critically impacts performance.

5.2 Compression Ratio Sensitivity

The bandwidth ratio $r = d_{msg}/d_{obs}$ directly controls message dimensionality. We ask: **to what extent can BVME preserve coordination effectiveness under aggressive compression, and how does this compare to deterministic compression (GACG)?**




Figure 6: On-path coupling is essential for effective bandwidth control. Learning curves comparing on-path (stochastic messages $z_i = \mu_i + \sigma_i \odot \epsilon$ fed to Q -network) versus off-path (deterministic mean $z_i = \mu_i$ fed to Q -network) at $r = 0.05$. On-path achieves 6-8% higher win rates by ensuring the KL penalty directly constrains decision-making representations.

Setup. We vary r on 3s5z from 0.02 \sim 0.30 and on 8m_vs_9m from 0.05 \sim 0.30, covering extreme bottlenecks (2–5%) to relaxed budgets (30%). For each setting, we measure the area under the learning curve (AUC) and final test win rate over 5 seeds. To visualize relative performance, we compute $\Delta AUC(r) = AUC_{BVME}(r) - AUC_{GACG}(r)$, where positive values indicate BVME achieves stronger cumulative learning than GACG at the same budget (Fig. 4). We further contrast BVME at low ratios ($r=0.05, 0.075, 0.10$) with GACG at full budget ($r=0.30$) in Fig. 5 and Tab. 1.

Results. Fig. 4 reveals a consistent U-shaped advantage curve. At very low r (≤ 0.05), BVME substantially outperforms GACG by preserving task-critical features through KL-regularized stochastic encoding: on 3s5z, $\Delta AUC = 6.0$ at $r = 0.02$ and 3.5 at $r = 0.05$; on 8m_vs_9m, gains reach 5.4 and 3.9 respectively. In the mid-range ($r \approx 0.10-0.15$), the gap narrows. On 8m_vs_9m, GACG even shows slight advantages ($\Delta AUC = -1.6$ at $r = 0.125, 0.15$), likely because deterministic projection suffices at these ratios while BVME’s KL penalty introduces regularization overhead. However, this regime is narrow: as r increases beyond 0.15, BVME recovers its advantage. At higher budgets ($r \geq 0.15$), BVME again outperforms GACG as KL regularization filters nuisance dimensions and prevents overfitting. On 3s5z, ΔAUC remains positive across all r values, indicating consistent benefits.

Extreme compression. Tab. 1 quantifies a striking result: BVME at $r = 0.05$ (83% message reduction vs. GACG-0.30) achieves *higher* win rates on both maps—0.916 vs. 0.868 on 3s5z and 0.841 vs. 0.788 on 8m_vs_9m—while also improving AUC. This demonstrates that variational encoding not only preserves coordination under severe bandwidth constraints but can actively improve performance by forcing compact, task-relevant representations.

These findings establish BVME’s robustness across bandwidth regimes: substantial gains under tight constraints, competitive performance in mid-ranges, and denoising benefits at high capacity. To validate the underlying mechanisms, we next isolate the role of on-path coupling, test alternative graph backbones, and sweep the regularization hyperparameters λ_{KL} and σ_0 .

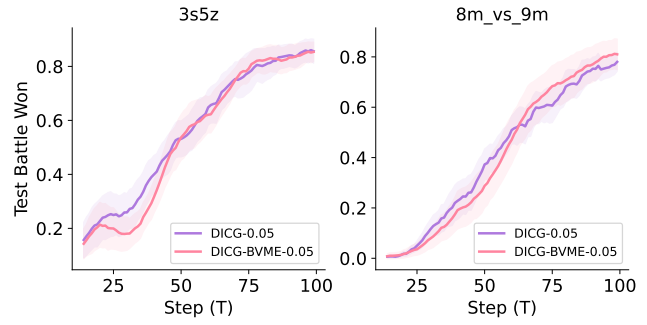


Figure 7: BVME provides modest gains on dense-graph backbones. Learning curves comparing DICG (fully connected) with and without BVME at $r = 0.05$. Final win rates are nearly identical on 3s5z (0.86 vs. 0.86) with only a 2% improvement on 8m_vs_9m (0.82 vs. 0.80), contrasting with 5-10% gains on sparse GACG.

5.3 On-Path vs. Off-Path Coupling

A core design choice in BVME is *on-path* coupling: using the sampled stochastic messages $z_i^{(l)}$ (Eq. 9) directly for Q -value estimation, ensuring the KL penalty constrains the representations that drive decisions. We validate this choice by comparing on-path training against an *off-path* variant that applies the same KL regularization but feeds only the deterministic mean $\mu_i^{(l)}$ into the Q -network.

Setup. Both variants use identical architectures and hyperparameters at bandwidth ratio $r=0.05$ on 3s5z and 8m_vs_9m, and deterministic messages ($z_i = \mu_i$) at evaluation for stability, differing only in how messages are used during training:

- **On-path:** Samples stochastic messages $z_i^{(l)} = \mu_i^{(l)} + \sigma_i^{(l)} \odot \epsilon$ and feeds them directly to the Q -network, ensuring the KL penalty constrains the representations used for decisions.
- **Off-path:** Computes the same KL divergence $KL(p_i^{(l)} || q)$ using $(\mu_i^{(l)}, \sigma_i^{(l)})$ but forwards only the deterministic mean $z_i^{(l)} = \mu_i^{(l)}$ for value estimation, decoupling regularization from the decision pathway.

Results. Fig. 6 shows on-path coupling achieves both faster convergence and higher asymptotic performance. On 3s5z, on-path reaches 0.88 win rate by 0.8M steps while off-path plateaus near 0.82; on 8m_vs_9m, the gap widens further with on-path achieving 0.80 vs. 0.72 for off-path. This gap arises because on-path training enforces consistency between the regularized representations and the representations used for control. The KL penalty directly shapes the stochastic messages that inform Q -values, creating tight coupling between compression and coordination. In contrast, off-path decouples this link: the encoder learns (μ_i, σ_i) under KL pressure, but since only μ_i drives decisions, the model can collapse $\sigma_i \rightarrow 0$ to minimize the penalty while encoding detailed information in μ_i —effectively bypassing the bandwidth constraint.

These results confirm that on-path coupling is essential for BVME’s effectiveness, ensuring the variational bottleneck meaningfully constrains the information used for coordination rather than just an auxiliary pathway.




Figure 8: Mean test win rates across KL weight λ_{KL} and prior scale σ_0 at $r=0.1$. Optimal settings differ by task: **3s5z peaks at ($\lambda_{\text{KL}}=0.5, \sigma_0=0.01$) achieving 0.927, while **8m_vs_9m** prefers ($\lambda_{\text{KL}}=2.0, \sigma_0=0.02$) achieving 0.882.**

5.4 Dense vs. Sparse Graph Backbones

To test whether BVME’s benefits generalize beyond sparse structures, we apply it to DICG—a fully connected backbone where all agents can communicate.

Setup. We compare DICG-0.05 against DICG-BVME-0.05 (with variational message encoding) at bandwidth ratio $r=0.05$ on 3s5z and 8m_vs_9m, using identical hyperparameters.

Results. Fig. 7 shows BVME provides only modest improvements on dense graphs: on 3s5z, final win rates are nearly identical (0.86 for both), while on 8m_vs_9m, BVME yields only 2% gain (0.82 vs. 0.80), contrasting sharply with 5-10% improvements on sparse GACG (Sec. 5.1). This aligns with our hypothesis: in fully connected topologies, agents aggregate complementary information from all teammates, providing rich information flow even under compression. BVME’s variational encoding offers limited benefit because agents need not rely on a few critical edges. Conversely, in sparse graphs, each edge carries more coordination responsibility, making message quality crucial. These results validate BVME’s design: it excels when communication channels are scarce, as KL-regularized encoding prioritizes task-critical information on limited edges, while gains are marginal on dense topologies where connectivity is already maximal.

5.5 Hyperparameter Sensitivity

BVME involves two key hyperparameters that jointly govern compression strength: the KL weight λ_{KL} controls penalty strength, while the prior variance σ_0^2 sets the baseline capacity. We investigate how their interaction affects coordination performance.

Setup. We perform a 3×3 grid search over $\lambda_{\text{KL}} \in \{0.5, 1.0, 2.0\}$ and $\sigma_0 \in \{0.005, 0.01, 0.02\}$, evaluating mean test win rates over 5 seeds on 3s5z and 8m_vs_9m, under fixed bandwidth ratio $r = 0.05$.

Results. Fig. 8 reveals task-dependent sensitivity to regularization strength. **3s5z:** Performance peaks at ($\lambda_{\text{KL}}=0.5, \sigma_0=0.01$) with 0.927 win rate, remaining strong across moderate settings but dropping to 0.868 at ($\lambda_{\text{KL}}=1.0, \sigma_0=0.01$), suggesting over-regularization. **8m_vs_9m:** The optimal setting shifts to ($\lambda_{\text{KL}}=2.0, \sigma_0=0.02$) achieving 0.882, with performance degrading substantially under weak regularization: ($\lambda_{\text{KL}}=0.5, \sigma_0=0.02$) yields only 0.783. BVME shows reasonable robustness (performance spans 6-10% across the grid) but benefits from task-specific tuning. **Moderate compression pressure extracts useful features, while extremes either fail**

to eliminate redundancy or remove critical information. The differing optimal settings across tasks suggest that coordination complexity determines the ideal regularization trade-off.

6 DISCUSSION

Computational Complexity. Let n denote the number of agents, d_{obs} the observation dimension, and $d_{\text{msg}} = r \cdot d_{\text{obs}}$ the message dimension with $r \ll 1$. BVME adds $\mathcal{O}(n \cdot d_{\text{msg}}^2)$ operations for two encoder MLPs ($\text{Enc}_\mu, \text{Enc}_\sigma$), $\mathcal{O}(n \cdot d_{\text{msg}})$ for reparameterized sampling, and $\mathcal{O}(n \cdot d_{\text{msg}})$ for closed-form KL divergence. Since $d_{\text{msg}}^2 = r^2 \cdot d_{\text{obs}}^2$ with $r \approx 0.05$, overhead is negligible compared to the baseline’s $\mathcal{O}(L \cdot n^2 \cdot d_{\text{obs}}^2)$ message-passing cost. Wall-clock training time increases by $< 5\%$ on SMAC; at evaluation, deterministic means ($z_i = \mu_i$) incur zero additional cost.

Relation to Prior Bandwidth-constrained Methods. MAGI [4] is the most closely related work; both apply information bottleneck principles [33] but with distinct objectives. MAGI bottlenecks **graph structure**, filtering noisy channels for adversarial robustness via dual-MI optimization. Conversely, BVME bottlenecks **individual messages**, regularizing Gaussian-sampled messages toward an uninformative prior to compress observations under hard constraints ($d_{\text{msg}} \ll d_{\text{obs}}$).

Key differences: **(1)** BVME provides explicit dimensional capacity control via compression ratio r and prior scale σ_0 with closed-form KL (Eq. 13), enabling practitioners to directly tune the bandwidth-performance trade-off; MAGI’s robustness terms regulate information flow but do not enforce dimensional constraints. **(2)** BVME uses lightweight diagonal-Gaussian encoding at a single GNN layer ($<5\%$ overhead), while MAGI requires nested variational bounds, categorical and Gaussian sampling across multiple layers, and adversarial training. **(3)** BVME’s on-path regularization ensures the KL penalty directly constrains representations used for decision-making (Sec. 5.3), rather than auxiliary pathways.

7 CONCLUSION

We introduced BVME, a variational message encoding framework for Bandwidth-constrained multi-agent coordination. Hard bandwidth constraints force selective encoding, but deterministic projections lack control over how compression occurs. BVME addresses this by modeling messages as samples from learned Gaussian posteriors regularized via KL divergence, providing tunable compression control through interpretable hyperparameters. BVME’s lightweight design—diagonal Gaussians, closed-form KL regularization, and on-path coupling—integrates seamlessly with existing coordination-graph methods while adding minimal overhead.

Across SMACv1, SMACv2, and MPE benchmarks, BVME achieves comparable or superior performance using 67–83% fewer message dimensions. Ablations reveal U-shaped bandwidth sensitivity, with BVME excelling at extreme compression ($r \leq 0.05$) and benefits most pronounced on sparse graphs where message quality critically impacts coordination. On-path coupling proves essential, substantially outperforming off-path regularization.

ACKNOWLEDGMENTS

This work is supported by the Australian Research Council under Australian Laureate Fellowships FL190100149.

REFERENCES

- [1] Stav Belogolovsky, Eran Iceland, Itay Naeh, Ariel Barel, and Shie Mannor. 2025. Interpretable Multi-Agent Communication via Information Gating. In *ICML 2025 Workshop on Collaborative and Federated Agentic Workflows*.
- [2] Wendelin Böhmer, Vitaly Kurin, and Shimon Whiteson. 2020. Deep Coordination Graphs. In *International Conference on Machine Learning (ICML)*, Vol. 119. PMLR, 980–991.
- [3] Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. 2019. TarMAC: Targeted Multi-Agent Communication. In *International Conference on Machine Learning (ICML)*. PMLR, 1538–1546.
- [4] Shifei Ding, Wei Du, Ling Ding, Jian Zhang, Lili Guo, and Bo An. 2024. Robust Multi-Agent Communication With Graph Information Bottleneck Optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 5 (2024), 3096–3107. <https://doi.org/10.1109/TPAMI.2023.3337534>
- [5] Wei Duan, Jie Lu, Yu Guang Wang, and Junyu Xuan. 2024. Layer-diverse Negative Sampling for Graph Neural Networks. *Transactions on Machine Learning Research* (2024).
- [6] Wei Duan, Jie Lu, and Junyu Xuan. 2024. Group-Aware Coordination Graph for Multi-Agent Reinforcement Learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, (IJCAI 2024), Jeju, South Korea, August 3-9, 2024*. 3926–3934.
- [7] Wei Duan, Jie Lu, and Junyu Xuan. 2025. Bayesian Ego-graph inference for Networked Multi-Agent Reinforcement Learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems (NIPS 2025)*.
- [8] Wei Duan, Jie Lu, and Junyu Xuan. 2025. Inferring Latent Temporal Sparse Coordination Graph for Multiagent Reinforcement Learning. *IEEE Transactions on Neural Networks and Learning Systems* 36, 8 (2025), 14358–14370. <https://doi.org/10.1109/TNNLS.2024.3513402>
- [9] Wei Duan, Junyu Xuan, Maoying Qiao, and Jie Lu. 2022. Learning from the Dark: Boosting Graph Convolutional Neural Networks with Diverse Negative Samples. In *Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022), Virtual Event*. AAAI Press, 6550–6558.
- [10] Wei Duan, Junyu Xuan, Maoying Qiao, and Jie Lu. 2024. Graph Convolutional Neural Networks With Diverse Negative Samples via Decomposed Determinant Point Processes. *IEEE Transactions on Neural Networks and Learning Systems* 35, 12 (2024), 18160–18171. <https://doi.org/10.1109/TNNLS.2023.3312307>
- [11] Benjamin Ellis, Jonathan Cook, Skander Moalla, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob N. Foerster, and Shimon Whiteson. 2023. SMACv2: An Improved Benchmark for Cooperative Multi-Agent Reinforcement Learning. In *The 36th Annual Conference on Neural Information Processing Systems (NeurIPS 2023), New Orleans, LA, USA, December 10–16*.
- [12] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 29.
- [13] Shengchao He, Hongzhi Ni, Jianhao Wang, Luo Wu, and Chongjie Zhang. 2024. Learning Multi-Agent Communication from Graph Modeling Perspective. In *International Conference on Learning Representations (ICLR)*.
- [14] Shariq Iqbal and Fei Sha. 2019. Actor-Attention-Critic for Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning (ICML)*. PMLR, 2961–2970.
- [15] Jiechuan Jiang, Chen Dun, Tiejun Huang, and Zongqing Lu. 2020. Graph Convolutional Reinforcement Learning. In *International Conference on Learning Representations (ICLR)*.
- [16] Jiechuan Jiang and Zongqing Lu. 2018. Learning Attentional Communication for Multi-Agent Cooperation. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 31.
- [17] Daewoo Kim, Sangwoo Moon, David Hostallero, Wan Ju Kang, Taeyoung Lee, Kyunghwan Son, and Yung Yi. 2019. Learning to Schedule Communication in Multi-agent Reinforcement Learning. In *7th International Conference on Learning Representations, ICLR 2019*.
- [18] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- [19] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- [20] Sheng Li, Jayesh K. Gupta, Peter Morales, Ross E. Allen, and Mykel J. Kochenderfer. 2021. Deep Implicit Coordination Graphs for Multi-agent Reinforcement Learning. In *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Virtual Event, United Kingdom*. ACM, 764–772.
- [21] Xiangyu Liu and Kaiqing Bai. 2023. Partially Observable Multi-agent RL with (Quasi-)Efficiency: The Blessing of Information Sharing. In *International Conference on Machine Learning (ICML)*. PMLR, 22106–22130.
- [22] Yong Liu, Weixun Wang, Yujing Hu, Jianye Hao, Xingguo Chen, and Yang Gao. 2020. Multi-Agent Game Abstraction via Graph Attention Neural Network. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020), New York, NY, USA*. AAAI Press, 7211–7218.
- [23] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *The 30th Annual Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA*. 6379–6390.
- [24] Hangyu Mao, Zhengchao Zhang, Zhen Xiao, Zhibo Gong, and Yan Ni. 2020. Learning Agent Communication under Limited Bandwidth by Message Pruning. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 5142–5149.
- [25] Frans A. Oliehoek and Christopher Amato. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer.
- [26] Thomy Phan, Fabian Ritz, Lenz Belzner, Philipp Altmann, Thomas Gabor, and Claudia Linnhoff-Popien. 2021. VAST: Value Function Factorization with Variable Agent Sub-Teams. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems (NIPS 2021), December 6-14, virtual*. 24018–24032.
- [27] Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. 2018. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018), Stockholm, Sweden, August 6-14, 2018*. PMLR, 4292–4301.
- [28] Mikayel Samvelyan, Tabish Rashid, Christian Schröder de Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob N. Foerster, and Shimon Whiteson. 2019. The StarCraft Multi-Agent Challenge. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2019), Montreal, QC, Canada*. International Foundation for Autonomous Agents and Multiagent Systems, 2186–2188.
- [29] Jianzhun Shao, Yao Lou, Hongchang Zhou, Shuncheng Jiang, and Xiangyang Ji. 2023. Complementary Attention for Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning (ICML)*. PMLR, 30780–30797.
- [30] Amanpreet Singh, Tushar Jain, and Sainbayar Sukhbaatar. 2019. Learning when to Communicate at Scale in Multiagent Cooperative and Competitive Tasks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- [31] Sainbayar Sukhbaatar, Rob Fergus, et al. 2016. Learning Multiagent Communication with Backpropagation. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 29.
- [32] Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Flores Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018), Stockholm, Sweden*. 2085–2087.
- [33] Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. The information bottleneck method. *arXiv preprint physics/0004057* (2000).
- [34] Guilherme S. Varela, Alberto Sardinha, and Francisco S. Melo. 2025. Networked Agents in the Dark: Team Value Learning under Partial Observability. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025, Detroit, MI, USA, May 19-23, 2025*. International Foundation for Autonomous Agents and Multiagent Systems / ACM, 2087–2095.
- [35] Anthony Wang, Songyuan Peng, Vijay Kumar, and Alejandro Ribeiro. 2024. Graph Neural Network-based Multi-agent Reinforcement Learning for Resilient Distributed Coordination of Multi-Robot Systems. In *IEEE International Conference on Robotics and Automation (ICRA)*.
- [36] Tonghan Wang, Jianhao Wang, Chongyi Zheng, and Chongjie Zhang. 2020. Learning Nearly Decomposable Value Functions Via Communication Minimization. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- [37] Jannis Weil, Zhenghua Bao, Osama Abboud, and Tobias Meuser. 2024. Towards Generalizability of Multi-Agent Reinforcement Learning in Graphs with Recurrent Message Passing. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, (AAMAS 2024), Auckland, New Zealand, May 6-10, 2024*. 1919–1927.
- [38] Qianlan Yang, Weijun Dong, Zhizhou Ren, Jianhao Wang, Tonghan Wang, and Chongjie Zhang. 2022. Self-Organized Polynomial-Time Coordination Graphs. In *International Conference on Machine Learning (ICML 2022), Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*. PMLR, 24963–24979.
- [39] Xiaoyu Yang, Jie Lu, and En Yu. 2025. Adapting Multi-modal Large Language Model to Concept Drift From Pre-training Onwards. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=b20VK2Gns5>
- [40] Xiaoyu Yang, Jie Lu, and En Yu. 2025. Walking the Tightrope: Autonomous Disentangling Beneficial and Detrimental Drifts in Non-Stationary Custom-Tuning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=1BAiQmAFsx>

- [41] Chao Yu, Akash Velu, Eugene Vinitzky, Jiaxuan Gao, Yu Wang, Alexandre M. Bayen, and Yi Wu. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [42] En Yu, Jie Lu, Kun Wang, Xiaoyu Yang, and Guangquan Zhang. 2025. Drift-aware collaborative assistance mixture of experts for heterogeneous multistream learning. *arXiv preprint arXiv:2508.01598* (2025).
- [43] En Yu, Jie Lu, Xiaoyu Yang, Guangquan Zhang, and Zhen Fang. 2025. Learning Robust Spectral Dynamics for Temporal Domain Generalization. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- [44] Ziluo Zhang, Tiejun Zhao, and Chongjie Meng. 2024. Multi-Agent Coordination via Multi-Level Communication. In *Advances in Neural Information Processing Systems (NeurIPS)*.