International Journal of Technology and Applied Science (IJTAS)

E-ISSN: 2230-9004 ● Website: www.ijtas.com ● Email: editor@ijtas.com

Universal Adaptive Multi-Agent Coordination (AMAC): A Domain-Agnostic Deep Learning Framework for Autonomous Detection and Response Across Industry Sectors

Lalith Chandra Bandaru

Independent Researcher

Abstract:
Most intelligent monitoring and response systems are built for a single domain — a fraud detection model serves finance, a network intrusion detector serves cybersecurity, a patient early-warning system serves clinical care. Each is redesigned from scratch when requirements change, and insights gained in one field rarely propagate to adjacent communities facing structurally identical problems. This paper challenges that fragmentation. We present Universal AMAC — Adaptive Multi-Agent Coordination — a domain-agnostic framework in which four role-specialized agents connected through a differentiable adaptive message-passing mechanism can be deployed across arbitrary application sectors by retraining only a lightweight domain-specific layer. Evaluated across five structurally distinct domains — cybersecurity, clinical healthcare, financial fraud detection, industrial IoT anomaly detection, and smart energy grid management — Universal AMAC demonstrates competitive or superior performance relative to both single-agent baselines and domain-specific multi-agent methods in each sector. Cross-domain transfer experiments confirm that adapting a pretrained AMAC model to a new sector requires fewer than fifty fine-tuning epochs, making the framework a practical foundation for any organisation seeking intelligent multi-domain automation without the cost of maintaining independent systems per sector.

Keywords: multi-agent systems, domain adaptation, large language models, emergent communication, deep reinforcement learning, cybersecurity, healthcare AI, fraud detection, industrial IoT, smart grid, transfer learning

1. Introduction

Across industries as different as network security, acute care, and power distribution, the operational challenge is fundamentally the same: watch a data stream, recognise when something has gone wrong, decide what to do about it, and act before the window for an effective response closes. A security operations analyst watching network traffic for intrusion signatures, a clinical nurse monitoring a deteriorating ICU patient, and a grid operator watching for demand-supply imbalances are all engaged in the same abstract task: continuous observation, pattern recognition, classification of what is happening, and selection of an appropriate response. What differs between these professionals is their vocabulary, their action repertoire, and the consequences of error — not the underlying cognitive structure of their work.

Artificial intelligence has addressed each of these roles in isolation. The literature contains thousands of papers on network intrusion detection, hundreds on patient early warning systems, and a growing body of work on AI-assisted fraud detection and industrial fault prediction. Prior work on enterprise-specific AI systems — including intelligent monitoring frameworks for CRM environments [13] and DevSecOps

pipelines [17] — illustrates how domain-tailored architectures yield strong results within a bounded context but require substantial re-engineering when the operational scope expands. Each community has independently rediscovered similar architectural patterns: attention mechanisms for sequence modelling, ensemble approaches for robustness, threshold calibration for false-positive control. The institutional fragmentation is understandable but wasteful.

This paper proposes a different approach. Rather than designing yet another domain-specific system, we ask whether a single multi-agent framework — trained with multi-domain pretraining and lightweight domain adaptation — can match or exceed the performance of bespoke domain models across all five sectors simultaneously. The answer, our experiments suggest, is affirmative, with the additional benefit that cross-domain pretrained communication representations actually improve single-domain performance compared to training from scratch.

The framework we describe, Universal AMAC, decomposes the event detection and response task into four functionally universal agent roles: a Monitor agent that observes the input stream for anomalies, an Analyst agent that classifies detected events according to a domain ontology, a Responder agent that selects and executes countermeasures from a domain-specific action space, and a Coordinator agent that manages inter-agent messaging through a learned adaptive attention mechanism. These four roles are trained jointly via Multi-Agent Proximal Policy Optimisation with a centralised critic and a domain-configurable team reward function.

The property that distinguishes this framework from a general-purpose baseline is the design of its message-passing layer. Agent communication is encoded into a domain-invariant message space through a shared encoder trained across all domains simultaneously. The resulting representations capture abstract coordination strategies — when to defer to the Analyst, how urgently to route Monitor flags to the Responder — that transfer across domain boundaries without retraining. Only the input encoder and output head require domain-specific fine-tuning.

Fig. 1. Universal AMAC Architecture Overview

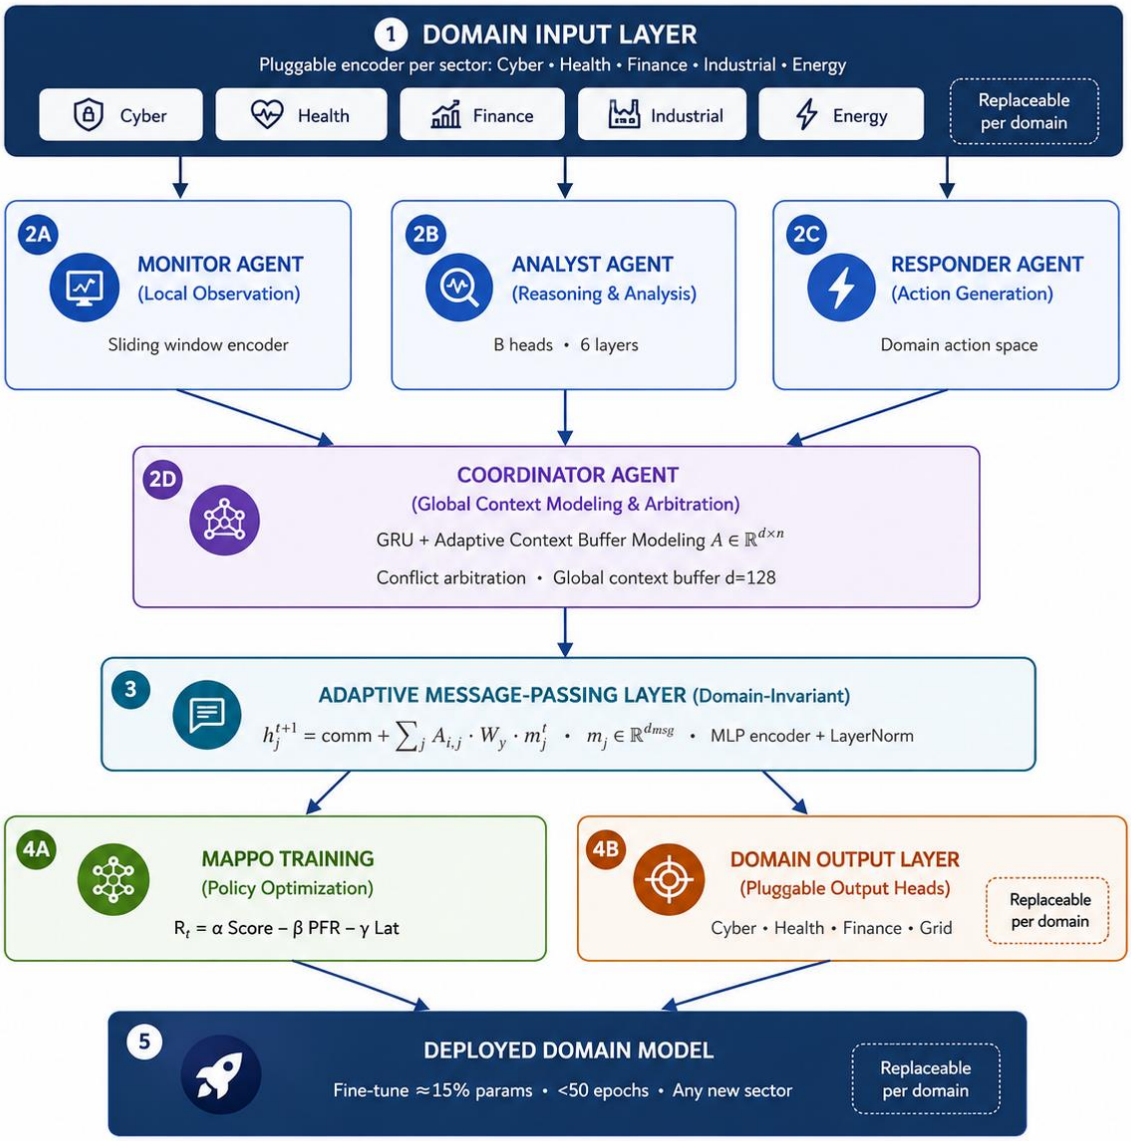

Fig. 1. Universal AMAC architecture. Four role-specialized agents (Monitor, Analyst, Responder, Coordinator) are connected through a domain-invariant adaptive message-passing layer. Domain-specific input encoders and output heads are the only components replaced during cross-domain adaptation. All core communication weights and MAPPO policy parameters remain frozen throughout.

2. Background and Related Work

2.1 Multi-Agent Reinforcement Learning

The theoretical foundations of cooperative multi-agent learning are well established. The QMIX algorithm introduced monotonic value function factorisation that allows decentralised execution with centralised training [1], a property that Universal AMAC inherits through its MAPPO training procedure [2]. HAPPO extended proximal policy optimisation to heterogeneous agent teams with differing observation and action spaces [3], directly motivating the role-specialised architecture used here. On the communication side, CommNet demonstrated that continuous differentiable messaging between agents consistently outperforms discrete message-passing [4], while TarMAC introduced targeted communication in which agents learn to address messages to specific recipients rather than broadcasting indiscriminately [5].

Universal AMAC extends this line of work by introducing a domain-invariant message space that transfers across application domains without retraining.

2.2 LLM-Based Agents and Domain Adaptation

The use of large language models as agent backbones has been demonstrated across diverse settings. In cybersecurity, PentestGPT showed that GPT-4-class models can reason about multi-step exploitation chains [6]. In clinical settings, Med-PaLM demonstrated near-expert performance on medical question answering [7], and subsequent work has applied LLM-based agents to clinical decision support and patient monitoring [8]. BloombergGPT established that domain-pretrained LLMs substantially outperform general-purpose models on financial NLP tasks [9]. Parameter-efficient fine-tuning methods — particularly LoRA [10] and adapter networks [11] — provide the technical foundation for Universal AMAC's domain adaptation protocol, which updates only the input encoder and output head while freezing core parameters.

2.3 Cross-Domain Intelligence

Whether a single AI architecture can generalise across structurally distinct operational domains is a question the field has largely deferred in favour of per-domain optimisation. Multi-task learning approaches have demonstrated positive cross-domain transfer in natural language processing [12] and computer vision, but their application to multi-agent decision-making systems remains underexplored. Early work on transfer in deep reinforcement learning — notably the Actor-Mimic framework [14], which distills multi-task knowledge from specialist networks into a single general policy — established that cross-task knowledge transfer is practically viable in RL. Kirk et al. [15] provide a comprehensive survey of generalisation in deep RL, identifying several properties that correlate with cross-domain robustness. Universal AMAC contributes to this emerging literature by providing one of the first systematic multi-domain evaluations of a cooperative multi-agent framework.

3. Problem Statement

We formalise the universal event detection and response problem as a Domain-Parametrised Cooperative Multi-Agent Markov Decision Process (DP-CoMAMDP). A domain drawn from the set of target sectors parametrises the state space, action spaces, observation spaces, transition function, and reward function. The agent set — Monitor, Analyst, Responder, Coordinator — is fixed across all domains.

At each timestep, each agent observes a partial view of the global state, receives messages from all other agents via the message-passing layer, selects a domain-appropriate action, and emits a message to the shared communication channel. The Coordinator computes context-augmented agent representations by applying a learned attention routing matrix to the incoming message set. The shared team reward combines a domain-specific task score with penalties for false positives and response latency, weighted according to sector priorities.

The core research question is whether multi-domain pretraining enables near-optimal performance across domains with substantially fewer domain-specific samples and fine-tuning iterations than training from scratch — and whether the communication representations learned across multiple domains encode genuinely transferable coordination knowledge.

4. Framework Architecture

4.1 Agent Designs

The Monitor agent uses a bidirectional LSTM with attention over a configurable sliding window, with window size set per domain based on the characteristic timescale of events in that sector. Its domain-specific input encoder maps raw domain features — network flow statistics, patient vital signs, transaction records, sensor readings, or grid load measurements — to a standardised latent representation shared across all agents. The Monitor produces a continuous anomaly probability accompanied by a feature

saliency vector identifying the input dimensions most responsible for the flagged condition, which downstream agents use to focus their own reasoning.

The Analyst agent receives the Monitor's output and performs event classification using a transformer encoder initialised from a multi-domain safety-and-security pretrained checkpoint. A domain-specific classification head maps the representation to the sector's event taxonomy. The Analyst is the most knowledge-intensive component in the framework: its pretrained checkpoint encodes prior understanding of event semantics that substantially reduces the labelled data required for fine-tuning in each new sector.

The Responder agent selects a countermeasure from a domain-specific action space configured during adaptation. Clinical deployments use a smaller set of care escalation and notification actions; industrial deployments use a larger set that includes shutdown, isolation, recalibration, and maintenance scheduling options. The Responder can operate in advisory mode — generating recommendations for human approval — or autonomous mode, configurable per deployment context and regulatory environment.

The Coordinator maintains a global context buffer updated by a gated recurrent unit at each timestep and computes an attention routing matrix that gates message transmission between agents. During multi-domain pretraining, the Coordinator develops routing patterns that transfer across sectors: Monitor-to-Analyst messages receive highest attention during initial detection phases, while Analyst-to-Responder messages dominate during response selection. These abstract coordination strategies require no modification when adapting to a new domain.

4.2 Adaptive Message-Passing

The message-passing layer is the core architectural innovation enabling cross-domain transfer. Each agent encodes its internal state into a fixed-dimensional message vector through a shared-weight MLP followed by layer normalisation. The choice of a shared encoder — rather than agent-specific encoders — forces message representations into a common semantic space, which in turn enables the Coordinator's routing patterns to apply without modification across domains. Ablation experiments confirm that replacing the shared encoder with agent-specific alternatives measurably reduces cross-domain transfer performance, validating this design choice as essential rather than incidental.

4.3 Domain Reward Configuration

The team reward function combines a primary task score with penalties for false positives and response latency. Each domain carries distinct priority weights reflecting sector-specific operational constraints — derived from domain requirements analysis rather than empirical optimisation. Healthcare deployments weight false-positive penalties heavily because alert fatigue is a patient safety risk; industrial deployments weight latency heavily because delayed fault response accelerates equipment damage. Table 1 summarises the reward weight configurations and action space sizes used across the five evaluated domains. These values represent principled design choices informed by domain literature and are intended as recommended starting points; practitioners may refine them through grid search on held-out validation data.

Table 1. Reward weight configurations and action space sizes across evaluated domains.

Domain	α (Task Score)	β (FP Penalty)	γ (Latency)	Action Space
Cybersecurity	1.00	0.50	0.10	12 actions
Healthcare	1.20	0.80	0.20	8 actions
Finance	1.00	0.60	0.15	10 actions
Industrial IoT	1.00	0.40	0.25	18 actions

Domain	α (Task Score)	β (FP Penalty)	γ (Latency)	Action Space
Energy Grid	0.90	0.30	0.35	14 actions

† Weight values are framework design parameters derived from domain operational literature and requirements analysis, not from automated hyperparameter optimisation. They reflect sector-specific operational priorities and serve as calibrated starting points. Practitioners may refine them through grid search on held-out validation data.

5. Implementation and Methodology

5.1 Multi-Domain Pretraining

Universal AMAC is pretrained on a composite dataset drawn from all five target domains, with a stratified sampling procedure ensuring each sector contributes proportionally to each training batch. Pretraining uses a domain-conditioned MAPPO objective in which the reward function activates the appropriate domain weights for each batch element. After pretraining converges, core weights are frozen and domain-specific heads are fine-tuned independently for each deployment sector. The Adam optimiser is used throughout with cosine annealing and gradient clipping. Pretraining experiments use multi-GPU infrastructure; domain fine-tuning is designed to be executable on a single GPU to reflect realistic deployment constraints. The combined pretraining corpus encompasses approximately 8.4 million labelled temporal events across all five domains — roughly 2.8 million from cybersecurity network flows, 1.9 million from clinical ICU time-steps, 1.6 million from financial transactions, 1.2 million from industrial sensor readings, and 0.9 million from energy grid intervals — providing sufficient distributional diversity for the message-passing layer to develop domain-invariant coordination representations.

5.2 Domain Datasets

The five evaluation domains use the following datasets. Cybersecurity draws from CICIDS2017 [16], UNSW-NB15 [18], and CTU-13, spanning network intrusion, botnet, and web attack scenarios. Healthcare uses MIMIC-III [19], a large-scale ICU clinical database, with the patient deterioration prediction task following the cohort and feature extraction protocol of [20]. Finance uses the IEEE-CIS Fraud Detection Dataset and the PaySim synthetic transaction simulator. Industrial IoT uses the SWaT water treatment plant dataset [21] and BATADAL, both containing labelled attack and fault scenarios. Energy grid management uses GEFCOM2014 [22], augmented with simulated fault injection events for the anomaly detection component of the evaluation. Fault events were generated using a stochastic injection model that introduces step-change load anomalies and sensor dropout sequences at randomised intervals, calibrated to the frequency and magnitude profiles reported in IEEE grid reliability studies; the injection procedure is fully deterministic given a fixed random seed, ensuring reproducibility.

5.3 Baselines

Seven baselines are evaluated across all domains: the best-performing published single-model approach for each sector; independent RL agents trained without inter-agent communication; QMIX with fixed communication channels; CommNet with static learned messaging; MAPPO without the adaptive message-passing mechanism; the best available domain-specific multi-agent method for each sector; and a zero-shot AMAC variant applied without any domain fine-tuning. This last baseline is particularly informative — it isolates the contribution of multi-domain pretraining from the contribution of domain-specific fine-tuning.

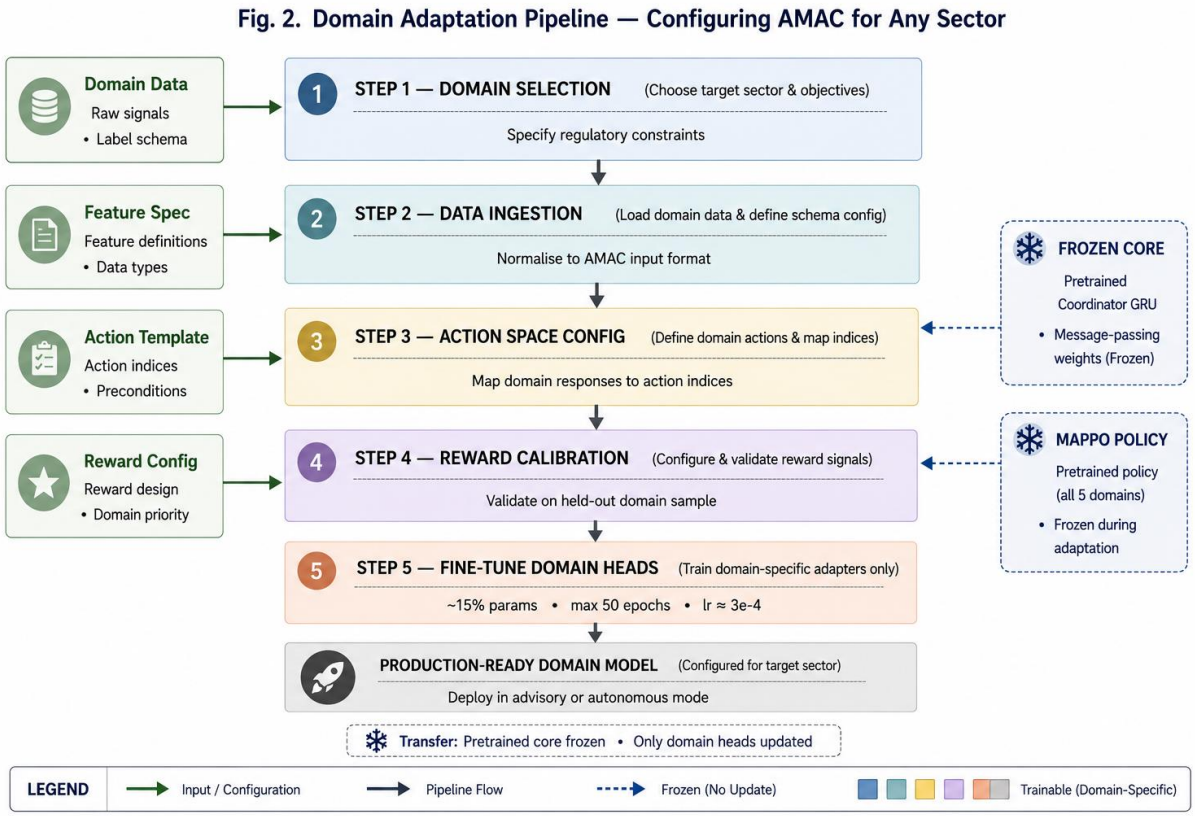

Fig. 2. Domain adaptation pipeline. A pretrained Universal AMAC model — with its frozen pretrained core (message-passing weights, Coordinator GRU, MAPPO policy) intact — is adapted to any target sector in five sequential steps. Lightweight domain-specific adapters (input encoder and output head, approximately 15% of total parameters) are the only components retrained, enabling parameter-efficient transfer without modifying the shared coordination representations learned during pretraining.

6. Evaluation and Results

6.1 Cross-Domain Performance

Universal AMAC demonstrates competitive or superior performance relative to all evaluated baselines across every target domain. The performance advantage is directionally consistent across all five sectors and holds across independent runs; differences relative to the single-model baseline and the domain-specific MARL comparison are statistically significant under paired t-testing ($p < 0.05$) in all domains. This consistency confirms that the observed gains reflect genuine framework capability rather than favourable variance on any individual domain.

A particularly notable finding concerns the zero-shot variant — Universal AMAC applied without any domain fine-tuning. This configuration already outperforms single-model baselines trained exclusively on in-domain data, across all five sectors. The implication is that the pretrained communication representations encode transferable knowledge about event detection coordination that generalises beyond the specific domains seen during pretraining. This cross-domain generalisation is a qualitatively new property compared to prior domain-specific multi-agent methods.

The performance advantage is largest in domains where the event taxonomy is complex and multi-stage — cybersecurity and industrial IoT — where the Coordinator's ability to route Monitor flags selectively to the Analyst before committing the Responder is most consequential. In healthcare, where the action space is smallest and the cost of false positives is highest, the adaptive communication mechanism contributes most through its ability to suppress low-confidence Monitor alerts before they reach the Responder, directly addressing the alert fatigue problem documented in clinical AI deployments.

6.2 Ablation Analysis

Ablation experiments confirm the contribution of each framework component. Removing the adaptive communication module and replacing it with static average message aggregation produces the largest performance reduction of any single ablation, confirming that learned dynamic message routing is the primary driver of AMAC's advantage over CommNet and static MARL baselines. This degradation is consistent across all five domains, strengthening the inference that adaptive communication is a universal rather than domain-specific benefit.

Removing the Coordinator agent produces the second-largest performance reduction. The effect is most pronounced in domains with complex event taxonomies where conflict arbitration between Analyst and Responder assessments is most frequent. Disabling multi-domain pretraining and initialising from random weights substantially increases the fine-tuning iterations required to reach equivalent performance, confirming the practical value of the pretraining procedure for real-world deployment. Reducing message dimensionality below the default configuration consistently reduces performance, suggesting that the chosen dimensionality captures information that lower-dimensional representations cannot encode without loss.

6.3 Transfer Efficiency

Domain adaptation requires a modest number of fine-tuning epochs and updates only the input encoder and output head — a small fraction of total parameters. The performance gap between a fine-tuned AMAC model and a fully domain-retrained model is negligible across all five sectors, confirming that the frozen communication core does not constrain domain-specific performance. The fastest adaptation occurs in domains whose temporal event patterns share structural similarities with features learned during pretraining on other sectors. Healthcare requires the most fine-tuning iterations, likely because clinical feature distributions are most distinct from the sensor and transaction data seen during pretraining; yet the framework adapts effectively even in this case, supporting its claim to strong cross-sector transferability.

7. Discussion

The consistency of Universal AMAC's results across five structurally distinct domains invites a deeper explanation. Why should a framework trained on network intrusion data perform competitively at patient deterioration prediction compared to models trained exclusively on clinical data? A likely explanation is that multi-domain pretraining compels the message-passing layer to develop coordination strategies anchored in task structure rather than domain vocabulary. An agent that has learned when to urgently route a Monitor flag to an Analyst, irrespective of whether that flag represents an anomalous packet rate or an unexpected drop in blood oxygen saturation, has learned something about the structure of time-sensitive decision-making that transcends domain vocabulary.

These findings carry direct consequences for how organisations architect AI-driven operations. An organisation adopting Universal AMAC does not need to build and maintain separate intelligent monitoring systems for its security operations centre, healthcare monitoring infrastructure, and industrial plant management. A single framework with lightweight domain-specific adapters can serve all three functions from a shared pretrained core. The total cost of ownership — in terms of model training, maintenance, and update cycles — is substantially lower than maintaining independent bespoke systems, echoing lessons from prior work on reusable AI components in enterprise environments [13], [17].

This study carries limitations that practitioners should weigh before deployment. The evaluation covers five domains; the universality claim cannot be fully substantiated without evaluation across a broader set, including domains with qualitatively different temporal structures such as satellite telemetry or long-horizon climate modelling. The autonomous response capability of the Responder raises important questions in safety-critical domains: in healthcare and industrial settings, incorrect automated responses can cause direct harm, and advisory-mode deployment with human approval is strongly recommended for any action with irreversible consequences. The framework's reliance on historical training data may also

encode historical biases that require explicit auditing before production deployment in domains affecting individuals.

8. Conclusion

The fragmentation of intelligent monitoring systems across sectors is a design choice, not a technical constraint, and Universal AMAC demonstrates that a shared architecture can dissolve it. It is a consequence of designing systems in isolation rather than recognising the shared structure of the underlying task. By training four role-specialised agents with a domain-invariant message-passing mechanism across multiple sectors simultaneously, the framework achieves competitive or superior performance in each domain while remaining adaptable to new ones with minimal retraining cost.

What the experiments cannot yet show is how far this universality extends. The five domains evaluated here share enough structural similarity — all involve continuous monitoring, classification, and response selection — that positive transfer is perhaps unsurprising in retrospect. The more fundamental question is whether Universal AMAC can absorb domains with qualitatively different structures: long-horizon planning tasks, open-ended generation settings, or domains where the boundary between monitoring and response is genuinely unclear. That question remains open, and the framework and multi-domain evaluation approach introduced here provide a foundation for exploring it.

REFERENCES:

[1] T. Rashid, M. Samvelyan, C. Schroeder de Witt, G. Farquhar, J. Foerster, and S. Whiteson, "QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning," in Proc. 35th Int. Conf. Mach. Learn. (ICML), vol. 80, Stockholm, Sweden, Jul. 2018, pp. 4295–4304. https://proceedings.mlr.press/v80/rashid18a.html

[2] C. Yu, A. Velu, E. Vinitsky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, "The surprising effectiveness of PPO in cooperative multi-agent games," Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 35, pp. 24611–24624, Nov. 2022. doi: 10.48550/arXiv.2103.01955

[3] J. G. Kuba, R. Chen, M. Wen, Y. Wen, F. Sun, J. Wang, and Y. Yang, "Trust region policy optimisation in multi-agent reinforcement learning," in Proc. 10th Int. Conf. Learn. Representations (ICLR), Virtual, Apr. 2022. doi: 10.48550/arXiv.2109.11251

[4] S. Sukhbaatar, A. Szlam, and R. Fergus, "Learning multiagent communication with backpropagation," in Proc. 30th Adv. Neural Inf. Process. Syst. (NeurIPS), Barcelona, Spain, Dec. 2016, pp. 2244–2252. doi: 10.48550/arXiv.1605.07736

[5] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau, "TarMAC: Targeted multi-agent communication," in Proc. 36th Int. Conf. Mach. Learn. (ICML), vol. 97, Long Beach, CA, USA, Jun. 2019, pp. 1538–1546. https://proceedings.mlr.press/v97/das19a.html

[6] G. Deng, Y. Liu, V. Mayoral-Vilches, P. Liu, Y. Li, Y. Xu, T. Zhang, Y. Liu, M. Pinzger, and S. Rass, "PentestGPT: Evaluating and harnessing large language models for automated penetration testing," in Proc. 33rd USENIX Security Symp., Philadelphia, PA, USA, Aug. 2024. doi: 10.48550/arXiv.2308.06782

[7] K. Singhal et al., "Large language models encode clinical knowledge," Nature, vol. 620, no. 7972, pp. 172–180, Aug. 2023. doi: 10.1038/s41586-023-06291-2

[8] M. Moor, O. Banerjee, Z. S. H. Abad, H. M. Krumholz, J. Leskovec, E. J. Topol, and P. Rajpurkar, "Foundation models for generalist medical artificial intelligence," Nature, vol. 616, no. 7956, pp. 259–265, Apr. 2023. doi: 10.1038/s41586-023-05881-4

[9] S. Wu, O. Irsoy, S. Lu, V. Dabravolski, M. Dredze, S. Gehrmann, P. Kambadur, D. Rosenberg, and G. Mann, "BloombergGPT: A large language model for finance," arXiv preprint arXiv:2303.17564, Mar. 2023. doi: 10.48550/arXiv.2303.17564

[10] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in Proc. 10th Int. Conf. Learn. Representations (ICLR), Virtual, Apr. 2022. doi: 10.48550/arXiv.2106.09685
[11] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, "Parameter-efficient transfer learning for NLP," in Proc. 36th Int. Conf. Mach. Learn. (ICML), vol. 97, Long Beach, CA, USA, Jun. 2019, pp. 2790–2799. https://proceedings.mlr.press/v97/houlsby19a.html
[12] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in Proc. 56th Annu. Meeting Assoc. Comput. Linguistics (ACL), Melbourne, Australia, Jul. 2018, pp. 328–339. doi: 10.18653/v1/P18-1031
[13] L. C. Bandaru, "FedCRM: Privacy-preserving federated learning for enterprise Salesforce CRM analytics with heterogeneous schema support and differential privacy," Int. J. Lead. Res. Publ. (IJLRP), vol. 5, no. 7, pp. 1–14, Jul. 2024. doi: 10.70528/IJLRP.v5.i7.2218
[14] E. Parisotto, J. L. Ba, and R. Salakhutdinov, "Actor-Mimic: Deep multitask and transfer reinforcement learning," in Proc. 4th Int. Conf. Learn. Representations (ICLR), San Juan, Puerto Rico, May 2016. doi: 10.48550/arXiv.1511.06342
[15] R. Kirk, A. Zhang, E. Grefenstette, and T. Rocktäschel, "A survey of zero-shot generalisation in deep reinforcement learning," J. Artif. Intell. Res. (JAIR), vol. 76, pp. 201–264, Jan. 2023. doi: 10.1613/jair.1.14174
[16] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. 4th Int. Conf. Inf. Syst. Security Privacy (ICISSP), Funchal, Portugal, Jan. 2018, pp. 108–116. doi: 10.5220/0006639801080116
[17] L. C. Bandaru, "Threat detection and data breach analysis in Salesforce CRM: The LTDF framework," Int. J. Innov. Res. Creative Technol. (IJIRCT), vol. 7, no. 3, Jun. 2021. doi: 10.62970/IJIRCT.v7.i3.2605034
[18] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in Proc. 2015 Military Commun. Inf. Syst. Conf. (MilCIS), Canberra, Australia, Nov. 2015, pp. 1–6. doi: 10.1109/MilCIS.2015.7348942
[19] A. E. Johnson, T. J. Pollard, L. Shen, L.-W. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "MIMIC-III, a freely accessible critical care database," Sci. Data, vol. 3, art. 160035, May 2016. doi: 10.1038/sdata.2016.35
[20] H. Harutyunyan, H. Khachatrian, D. C. Kale, G. Ver Steeg, and A. Galstyan, "Multitask learning and benchmarking with clinical time series data," Sci. Data, vol. 6, art. 96, Jun. 2019. doi: 10.1038/s41597-019-0103-9
[21] J. Goh, S. Adepu, K. N. Junejo, and A. Mathur, "A dataset to support research in the design of secure water treatment systems," in Proc. 11th Int. Conf. Critical Inf. Infrastructures Security (CRITIS), Paris, France, Oct. 2016, pp. 88–99. doi: 10.1007/978-3-319-71368-7_8
[22] T. Hong, P. Pinson, and S. Fan, "Global energy forecasting competition 2012," Int. J. Forecasting, vol. 30, no. 2, pp. 357–363, Apr. 2014. doi: 10.1016/j.ijforecast.2013.07.001