Local Anomaly Detection with Partial Observation in Multi-agent Systems as a Data Matching Game

Extended Abstract

Zixin Ye
University of Melbourne
Parkville, Australia
zixin4@student.unimelb.edu.au

Tansu Alpcan
University of Melbourne
Parkville, Australia
tansualpcan@gmail.com

Christopher Leckie
University of Melbourne
Parkville, Australia
caleckie@unimelb.edu.au

ABSTRACT

Local anomaly detection in a multi-agent system is a pervasive but challenging problem. The challenge entails how agents with heterogeneous objectives and partial data collection train local anomaly detectors for heterogeneous domain-specific tasks. This paper proposes a distributed training method to address this question. Our approach involves a game-theoretic framework to address agents' heterogeneous objectives and a transformer-based model to handle partial data observation. Our game, conditionally proven as a potential game, guides agents under the same local objectives into a data-sharing group for local training. Compared to other top-performing SOTAs, our evaluation outcomes empirically reflect the efficiency and robustness of our method in multi-agent scenarios.

KEYWORDS

Anomaly Detection, Game Theory, Multi-agent Systems

ACM Reference Format:
Zixin Ye, Tansu Alpcan, and Christopher Leckie. 2025. Local Anomaly Detection with Partial Observation in Multi-agent Systems as a Data Matching Game: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Agent-level anomaly detection is vital in multi-agent systems, such as the Internet of Things [2, 6] and Cognitive Radio Networks [8, 9]. Unlike many central solutions [5, 13], agent-level operations lack full-system observations and come with heterogeneous anomaly detection objectives. Examples include product quality control from multiple suppliers. Suppliers providing different materials assess product quality with local expertise and partial observations, which convert the anomaly detection problem into a multi-agent problem with local objectives and data access.

Data-sharing is the biggest challenge in multi-agent anomaly detection problems, especially under the lack of central control. As shown in Fig.1, without accessing other agents' information, what is one agent's optimal data-sharing strategy that 1) identifies other agents under the same local objectives and 2) shares data to train a robust anomaly detector with limited observations? Many previous works assume agents' anomaly detection tasks have a singular objective with complete system observation, including contextual [2, 7, 12, 16], and multi-view [11, 14] anomaly detection methods. Without such a central control assumption, the anomaly detection model may not know which local context an agent is situated in, thus being unable to group correct contextual information from agents to each local anomaly detector.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

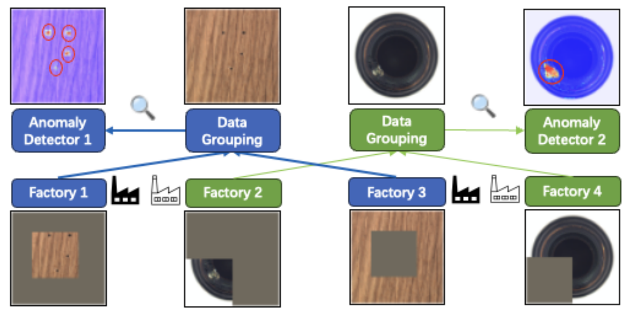

Figure 1: Multiple factories (as agents) share local data for aggregated anomaly detection. Factories producing the same product types have common anomaly detection objectives, whose data sharing provides more context to each other.

This paper proposes a novel multi-agent anomaly detection method that tackles the above-mentioned challenges. Our method solves two significant challenges introduced by multi-agent settings: 1) how to create robust anomaly detection against partially observable input data and 2) how to guide agents from the same local anomaly detection problem to share information and solve their local problems simultaneously. Our first innovation is training multiple local anomaly detectors robustly against agents' imperfect data contributions. We applied the masked auto-encoding transformer as our anomaly detectors' backbone, a.k.a. MAETAD. Our second innovation is to guide agents to contribute local data toward their most relevant local problem-solving anomaly detectors. We proposed a non-cooperative data-matching game where each agent selects and contributes data to a pre-trained local anomaly detector. This paper provides an overview of our methods.

2 BACKGROUND

A well-known communication structure [3] in multi-agent systems is formulated as a bipartite graph $\{\mathcal{V}, C, \mathcal{A}\}$. $C := \{c_1, \cdots, c_M\}$ is the set of aggregation nodes grouping input data for local anomaly detectors. $\mathcal{V} := \{v_1, \cdots, v_N\}$ is the set of agents requesting local anomaly detection on their collected data, and $\mathcal{A} := \{A \in$

$\{0,1\}^{N\times M}|\mathbf{1}^T\cdot A=\mathbf{1}\}$ is the agents' connection matrix from $\mathcal{V}$ to C. Data collected by each agent is denoted as a set of n_i vectors $D_{v_i}\in\mathbf{R}^{n_i\times d}$. Anomaly detector aggregates connected agents' data into an input matrix $D_{c_k}\in\mathbf{R}^{L\times d}$, where the position of each D_{v_i} are determined by an agent-reported positioning matrix, $E_{v_i}\in\{0,1\}^{n_i\times L}$. Data matching between agents and anomaly detectors can be expressed as $D_{v_i}=E_{v_i}D_{c_k}$ and $D_{c_k}=\sum_{v_i\in c_k}E_{v_i}^TD_{v_i},\forall v_i\in c_k$.

3 PROBLEM STATEMENT

We consider a multi-agent system with N agents sparsely distributed over a large area that provides local data to M local anomaly detection problems. Anomaly detectors may have partially observable input data due to limited agents' contributions, and agents do not know which anomaly detector is trained to solve their local problems. (1) How do we train each local anomaly detector to be robust against partially observable input data during decision time? (2) How do agents find and send their data to the correct local anomaly detectors that solve their local problems?

4 MAETAD: A NOVEL LOCAL ANOMALY DETECTOR

As our solution to Problem 1, we proposed a novel anomaly detector, MAETAD, to solve local anomaly detection problems. MAETAD is realized by a masked-autoencoding transformer structure to improve robustness against agents' partial data contributions in $D_{c_k}\in\mathbf{R}^{L\times d}$. The MAETAD model consists of an encoder and decoder concatenated as the function $f_{c_k}(\cdot,\cdot;\theta_f):(\mathbf{R}^{L\times d},\mathbf{R}^L)\rightarrow\mathbf{R}^{L\times d}$. Mathematically, given input data $D_{c_k}\in\mathbf{R}^{L\times d}$ partially observable in the L positions in E_{c_k}, the model output $f_{c_k}(D_{c_k},E_{c_k};\theta_f)\in\mathbf{R}^{L\times d}$ reconstructs all L-positioned data in D_{c_k} in the model output. The random masking simulates partially observable data in the model input, and the ability to handle random masked positions empowers MAETAD to detect local anomalies from limited input observations. Our training loss function is the anomaly detection version of the Hyper-sphere Classifier defined in [2] as follows,

$$l_{c_k}(D_{c_k},E_{c_k};\mathbf{y},\theta_f)=\sum_{j=1}^{L}(1-\mathbf{y}_j)||(f_{c_k}(D_{c_k},E_{c_k};\theta_f)-D_{c_k})_j||_2^2 - \mathbf{y}_j\log(1-e^{-||(f_{c_k}(D_{c_k},E_{c_k};\theta_f)-D_{c_k})_j||_2^2}), \tag{1}$$

where $||(f_{c_k}(D_{c_k},E_{c_k};\theta_f)-D_{c_k})_j||_2^2$ is the mean-square loss of data at position j, and the binary vector $\mathbf{y}\in\{0,1\}^L$ represents the ground-truth anomaly labels at all L positions of the input data. In our one-class-learning setting, only the first term is preserved as $\mathbf{y}_l=0,\forall l=0,\cdots,L$, which eliminates the second term.

As our solution to Problem 2, we rigorously analyze our proposed multi-agent anomaly detection task in a game-theoretic framework. We define our game as a tuple: $\mathcal{G}=\{\mathcal{V},\mathcal{A},\mathcal{U}\}$. The player set consists of all agents in $\mathcal{V}$, whose strategy profile is presented as the connection matrix $A\in\mathcal{A}$, where A_i is the device's v_i action. The last component $\mathcal{U}$ represents the space of utility functions for all local agents U_{v_i}. Component $\mathcal{U}$ contains the set of the agents' utility functions, denoted as $\{u_{v_i}:\mathcal{A}\rightarrow\mathcal{R}|v_i\in\mathcal{V}\}$. We first express the utility given by the aggregation node c_k as

$$u_{v_i}(A_i;A_{-i})=-\sum_{k=1}^{M}A_i^k||E_{v_i}f_{c_k}(\sum_{i=1}^{N}A_i^kE_{v_i}^TD_{v_i})-A_i^kD_{v_i}||_2^2, \tag{2}$$

The best response for each player v_i is the action A_i that maximizes (2). With agents searching for their best responses simultaneously, their optimal solutions route local data D_{v_i} to the aggregation nodes under the same local anomaly detection problems. We formulate the distributed algorithm Alg.(1) to depict the procedures of such data-matching, equivalently, the realization of our best response dynamic.

Algorithm 1: Data Matching of v_i

Input: Network Parameters A_{-i}, $\{D_{v_i}\}_{v_i\in\mathcal{V}}$, $\{f_{c_k}\}_{c_k\in C}$
Output: Connection Strategy A_i
function $\arg\max_{A_i} u_{v_i}(A_{-i},\{D_{v_i}\}_{v_i\in\mathcal{V}},\{f_{c_k}\}_{c_k\in C})$
  Step 1: Return the A_i that maximizes (2)
  Step 2: Update connections to f_{c_k} with the new A_i

5 EXPERIMENT RESULTS

We conduct our experiments on the anomaly detection datasets in MVTec-AD [1], a manufacturing quality control benchmark containing 10 objects and 5 textures products. The compared anomaly detection models are three top-performing state-of-the-art algorithms [4, 10, 15] on MVTec-AD. As shown in Table 1, we first examine the anomaly detection performance between MAETAD and three compared models in the presence of randomly missing input values. Then, we simulate and visualize the best response dynamics of our proposed game. Finally, we carried out the ablation studies to demonstrate the attributions of our game-theoretic model training methods. Our distributed model training empirically surpasses the central state-of-the-art algorithms by 5% in AUROC and 17% in AUPR on various benchmark datasets.

Metrics	**AUPR(%)**		**AUROC(%)**	
Products	**textures**	**objects**	**textures**	**objects**
FF [15]	31.4	41.7	77.4	65.1
PADIM [4]	24.1	22.8	73.2	74.1
PC [10]	25.1	28.4	77.2	80.4
MAETAD (OURS)	**49.5**	**41.7**	**80.6**	**86.1**

Table 1: Evaluations are performed in Averaged AUPR/AUROC (%) on two product categories in MVTec-AD datasets. Our MAETAD model outperforms three SOTAs in both textures' and objects' anomaly detection tasks

ACKNOWLEDGMENTS

This work was supported in part by the Australian Research Council Linkage Project under the grant LP190101287, and by Northrop Grumman Mission Systems' University Research Program. The authors would like to thank Justin Kopacz, Kerry Brown, and Tarun Soni of Northrop Grumman Corporation for providing valuable insights and feedback throughout this research.

REFERENCES

[1] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. 2019. MVTec AD — A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 9584–9592. https://doi.org/10.1109/CVPR.2019.00982

[2] Chris U. Carmona, Franois-Xavier Aubet, Valentin Flunkert, and Jan Gasthaus. 2022. Neural Contextual Anomaly Detection for Time Series. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 2843–2851. https://doi.org/10.24963/ijcai.2022/394 Main Track.

[3] K.-C. Chen, Y.-J. Peng, N. Prasad, Y.-C. Liang, and S. Sun. 2008. Cognitive Radio Network Architecture: Part II – Trusted Network Layer Structure. In *Proceedings of the 2nd International Conference on Ubiquitous Information Management and Communication* (Suwon, Korea) *(ICUIMC '08)*. Association for Computing Machinery, New York, NY, USA, 120–124. https://doi.org/10.1145/1352793.1352818

[4] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. 2021. PaDiM: A Patch Distribution Modeling Framework for Anomaly Detection and Localization. In *Pattern Recognition. ICPR International Workshops and Challenges*, Alberto Del Bimbo, Rita Cucchiara, Stan Sclaroff, Giovanni Maria Farinella, Tao Mei, Marco Bertini, Hugo Jair Escalante, and Roberto Vezzani (Eds.). Springer International Publishing, Cham, 475–489.

[5] Keju Huang, Jun Yang, Hui Liu, and Pengjiang Hu. 2022. Deep Learning of Radio Frequency Fingerprints from Limited Samples by Masked Autoencoding. *IEEE Wireless Communications Letters* (2022), 1–1. https://doi.org/10.1109/LWC.2022.3184674

[6] Sheo Yon Jhin, Jaehoon Lee, and Noseong Park. 2023. Precursor-of-Anomaly Detection for Irregular Time Series. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. Association for Computing Machinery, New York, NY, USA, 917–929. https://doi.org/10.1145/3580305.3599469

[7] Seongwoo Kim, He Cai, Cunqing Hua, Pengwenlong Gu, Wenchao Xu, and Jeonghyeok Park. 2020. Collaborative Anomaly Detection for Internet of Things based on Federated Learning. In *2020 IEEE/CIC International Conference on Communications in China (ICCC)*. 623–628. https://doi.org/10.1109/ICCC49849.2020.9238913

[8] Erma Perenda, Sreeraj Rajendran, Gerome Bovet, Sofie Pollin, and Mariya Zheleva. 2021. Learning the Unknown: Improving Modulation Classification Performance in Unseen Scenarios. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications* (Vancouver, BC, Canada). IEEE Press, 1–10. https://doi.org/10.1109/INFOCOM42981.2021.9488835

[9] Sreeraj Rajendran, Wannes Meert, Vincent Lenders, and Sofie Pollin. 2019. Unsupervised Wireless Spectrum Anomaly Detection With Interpretable Features. *IEEE Transactions on Cognitive Communications and Networking* 5, 3 (2019), 637–647. https://doi.org/10.1109/TCCN.2019.2911524

[10] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. 2022. Towards Total Recall in Industrial Anomaly Detection. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 14298–14308. https://doi.org/10.1109/CVPR52688.2022.01392

[11] Xiang-Rong Sheng, De-Chuan Zhan, Su Lu, and Yuan Jiang. 2019. Multi-View Anomaly Detection: Neighborhood in Locality Matters. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 4894–4901. https://doi.org/10.1609/aaai.v33i01.33014894

[12] Xiuyao Song, Mingxi Wu, Christopher Jermaine, and Sanjay Ranka. 2007. Conditional Anomaly Detection. *IEEE Transactions on Knowledge and Data Engineering* 19, 5 (2007), 631–645. https://doi.org/10.1109/TKDE.2007.1009

[13] Yves Teganya and Daniel Romero. 2022. Deep Completion Autoencoders for Radio Map Estimation. *IEEE Transactions on Wireless Communications* 21, 3 (2022), 1710–1724. https://doi.org/10.1109/TWC.2021.3106154

[14] Zhen Wang and Chao Lan. 2020. Towards a Hierarchical Bayesian Model of Multi-View Anomaly Detection. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, Christian Bessiere (Ed.). International Joint Conferences on Artificial Intelligence Organization, 2420–2426. https://doi.org/10.24963/ijcai.2020/335 Main track.

[15] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. 2021. FastFlow: Unsupervised Anomaly Detection and Localization via 2D Normalizing Flows. arXiv:2111.07677 [cs.CV]

[16] Xiang Yu, Hui Lu, Xianfei Yang, Ying Chen, Haifeng Song, Jianhua Li, and Wei Shi. 2020. An adaptive method based on contextual anomaly detection in Internet of Things through wireless sensor networks. *International Journal of Distributed Sensor Networks* 16, 5 (2020), 1550147720920478. https://doi.org/10.1177/1550147720920478 arXiv:https://doi.org/10.1177/1550147720920478