Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Cross-Domain Knowledge Transfer Algorithms For Rapid Agent Adaptation In New Environments

Dr. Shanthi Vairavan [1]*, K. Priyadharshini [2], M. Selin3, Dr.P. Balamurugan[4]

[1]Professor & Principal, Department of Computer Science, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. Email: shanthiv@maher.ac.in

2Assistant Professor, Department of Management Studies, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. Email: kpriyamba@maher.ac.in

3Department of Computer Applications, MES College Marampally, Ernakulam, Kerala, India. selin.m.a@gmail.com

4Associate Professor, Department of Networking and Communications, SRM Institute of Science and Technology, Kattankulathur, Tamil Nadu, India. E-mail: balamurp@srmist.edu.in

*Corresponding author: Email: shanthiv@maher.ac.in

Abstract

A major challenge in AI and machine learning is the ability to quickly adapt an intelligent agent to a new, and often very different, environment. In dynamic environments, conventional reinforcement learning and supervised learning methods tend to rely on a lot of task-specific data and lengthy training periods, which can hinder efficiency. In this paper, we propose a novel cross-domain knowledge transfer framework that enhances the adaptation speed of agents by combining domain-invariant feature extraction, the reuse of policies, and structured knowledge distillation. The framework allows for the effective transfer of knowledge from past experiences, while also reducing "negative transfer" of information as the agents learn in a new environment. The experiments were run on the modified version of LunarLander, CartPole, and MountainCar tasks, along with the image-based cross-domain tasks using MNIST and Fashion-MNIST datasets. The following measures were used to evaluate: cumulative reward, task success rate, adaptation speed, and computational efficiency. The proposed framework demonstrated a consistent improvement over the standard reinforcement learning, fine-tuning, and MAML-based adaptation methods, with the average task success rate of 92%, cumulative rewards of up to 220, and a decreased number of adaptation episodes by over 50%. The ablation studies further demonstrated that all the components are important and make a large impact on improving learning efficiency, especially domain-invariant feature extraction. The outcomes demonstrated the framework's ability to transfer knowledge to a variety of environments, suggesting its potential for use in robotics, autonomous systems, and multi-agent coordination. Going forward, the work will include scaling to highly divergent domains, adaptive weighting strategies, and real-world deployments to assess robustness and applicability further.

Keywords: Cross-domain learning, Knowledge transfer, Agent adaptation, Policy reuse, Domain-invariant features, Reinforcement learning, multi-agent systems.

1. Introduction

Adapting intelligent agents to new and unfamiliar environments is a key challenge in the field of artificial intelligence and machine learning [2]. Traditional learning approaches, such as reinforcement learning and supervised learning, require a lot of task-specific training data and interaction with the environment, making them inefficient in dynamic or unseen environments [1][13]. A potential solution is cross-domain knowledge transfer, where agents can utilize knowledge that they have gained in previous tasks or environments to help them learn faster in a new environment [4][5].

This paper tackles the challenge of creating algorithms to enable the effective and efficient transfer of knowledge across heterogeneous domains [11]. The emphasis is on empowering agents to quickly adjust to new environmental conditions in order to perform at an optimum level with minimum training time [9]. Some important goals are to create tools for: Domain-invariant Feature Extraction, Policy Reuse and Structured

Knowledge Distillation, so that prior learning can be used to aid new tasks without causing negative transfer to them [10][17].

The study also highlights the need to measure adaptation performance in terms of common indicators like task success rate, accumulated reward, and adaptation speed. This section aims to point out that transfer across domains is neither easy nor trivial, and outlines the motivation for creating strong algorithms that increase the agent's ability to be flexible and generalizable. The proposed approach seeks to develop scalable and practical solutions that can be applied to a variety of AI applications, such as robotics, autonomous systems, and multi-agent environments.

Key Contributions

- Creation of a novel cross-domain knowledge transfer framework to help agents easily adjust to new and diverse environments.
- New feature representations and policies are introduced to ensure domain-invariant representation and reusing features to reduce negative transfer and enhance adaptation efficiency.
- Detailed comparison against benchmark datasets and simulations, with significant increases in task success rates, faster convergence, and better generalization than traditional learning approaches.

In Section I, the research problem and the necessity of quick cross-domain adaptation are introduced, along with the research goals of this study. The existing studies on transfer learning, domain adaptation, and multi-agent learning are summarized, and research gaps are identified in Section II. In Section III, the proposed methodology is introduced, which involves algorithmic design, domain mapping, and knowledge distillation mechanisms. The experimental setup, datasets, evaluation metrics, and hardware/software configurations are explained in Section IV. The results, performance analysis, and comparative evaluation are discussed in Section V, and the conclusion and future directions are discussed in Section VI.

2. Literature Survey

Cross-domain knowledge transfer has become a key technique in AI, crucial to fast adaptation in new environments [3][7]. The existing research efforts in transfer learning have concentrated on improving the learning effectiveness by using pre-trained models to lessen the need for an adequate amount of data in the target task [6]. Fine-tuning and feature-based transfer have been successful in homogeneous domains, but have been unsuccessful when applied to heterogeneous or vastly different domains. The main purpose of domain adaptation is to overcome this limitation by trying to match feature distributions from the source to the target domain, but it may suffer from scalability and generalization problems in highly variable target domains.

In recent years, meta-learning techniques, such as model-agnostic meta-learning (MAML) and variants, have gained significant traction for enabling agents to acquire new tasks with low training costs [8][14]. They focus on fast adaptation by fine-tuning model parameters to quickly adjust to the task, but they may impose high computational costs, and the requirement for task similarity may be a limiting factor for practical use [12][18]. There has been a great deal of research on multi-agent learning systems that have investigated learning through collaboration to speed up adaptation, while ensuring that the agents do not interfere with each other or that there is no negative transfer [15][16].

More recently, knowledge distillation and policy reuse have been shown to be effective ways to efficiently transfer experience across tasks. These techniques allow the critical data to be preserved and reduce the chance of degradation of performance when using them in new domains. While these developments have been made, there is still a lack of scalable algorithms for integrating domain-invariant feature extraction, policy reuse, and structured knowledge distillation to obtain good adaptation results across dynamic and diverse environments[19][20]. In this study, we try to tackle this by proposing a comprehensive cross-domain knowledge transfer framework for fast and efficient agent adaptation.

3. Methodology

The proposed approach is directed towards the possibility of fast transfer of knowledge from one domain to another in order to adapt agents to new environments. The framework aims to build on prior learning and optimize negative transfer and learning efficiency in disparate domains. Structured Knowledge Distillation, Policy Reuse, and Domain Invariant Feature Extraction are the key elements of the methodology.

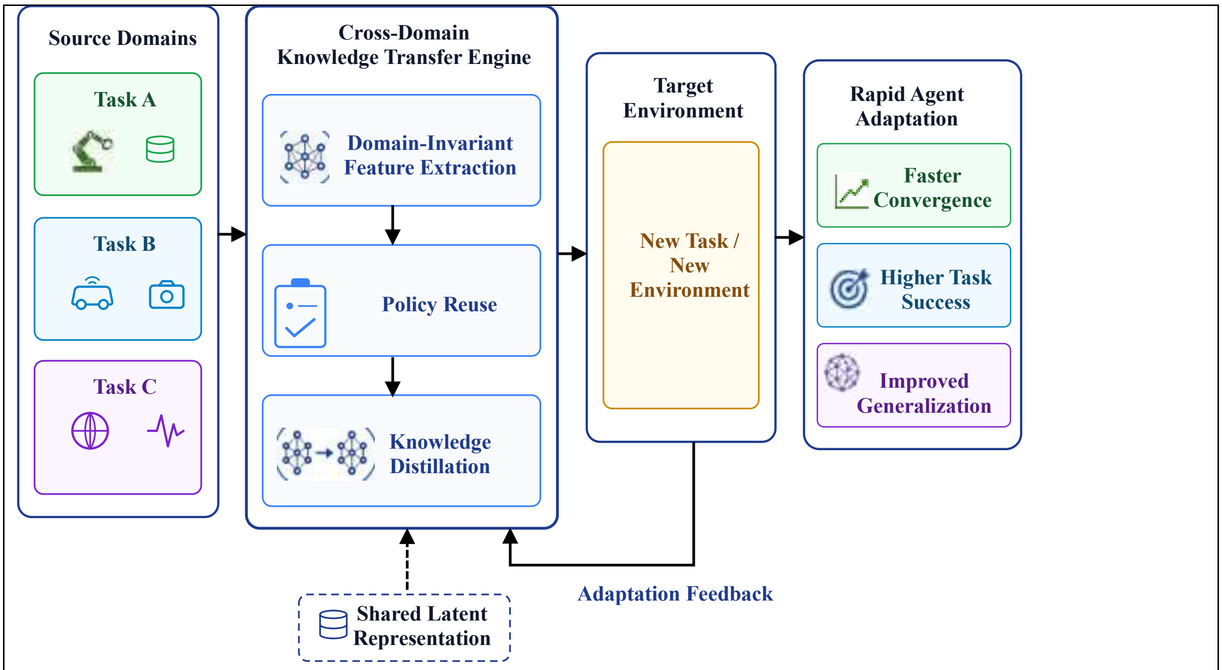

Figure 1: Conceptual framework for cross-domain knowledge transfer for rapid agent adaptation

The cross-domain knowledge transfer framework is proposed in Figure 1. In general, the central Knowledge Transfer Engine integrates domain-invariant feature extraction, policy reuse, and knowledge distillation techniques, which are fed from source domains with a variety of tasks. Knowledge Consolidation through a shared latent representation. The engine produces adaptation strategies for the target environment, which allow the agent to rapidly adapt and converge with a high success rate and generalization. The engine is continuously improved as a result of data from the target environment.

Domain-Invariant Feature Extraction

This component makes sure features gained from the source domain are relevant and transferable to the target domain. Each domain has a neural feature encoder that transforms the high-dimensional observation into a shared latent space, thereby mitigating domain-specific biases. Formally, given source domain data $X_s$and target domain data X_t, the encoder $f_\theta$produces representations $Z_s = f_\theta(X_s)$and $Z_t = f_\theta(X_t)$that minimize the distribution divergence $D(Z_s, Z_t)$, such as Maximum Mean Discrepancy (MMD):

$$\mathcal{L}_{\text{MMD}} = \| \frac{1}{n_s}\sum_{i=1}^{n_s} \phi(Z_s^i) - \frac{1}{n_t}\sum_{j=1}^{n_t} \phi(Z_t^j) \|^2 \qquad (1)$$

Equation (1) shows the MMD-based domain alignment loss

Policy Reuse Mechanism

Existing action-value functions are weighted averages and used to adapt previously learned policies from source tasks to the target environment. Let $Q_s$denote the source domain Q-function and $Q_t$the target domain Q-function. The updated policy $Q_t'$is calculated as:

$$Q_t'(s,a) = \alpha Q_s(s,a) + (1-\alpha) Q_t(s,a) \qquad (2)$$

Where in equation (2) $\alpha \in [0,1]$controls the contribution of prior knowledge, allowing agents to balance between exploration and exploitation in the new environment.

Structured Knowledge Distillation

It is a technique that passes on the important information from source models to target models in a condensed format. The soft targets created from source models help the learning of target models, saving time and retaining essential behaviors. The distillation loss $\mathcal{L}_{KD}$ is defined as:

$$\mathcal{L}_{KD} = -\sum_i P_s^i \log P_t^i \qquad (3)$$

where in equation (3) $P_s$and P_t are the softmax outputs of the source and target models, respectively.

The proposed methodology combines these elements to allow agents to effectively use their previous knowledge, adapt quickly to new environments, and perform well in various domains. The framework can be applied to various AI applications, such as robotic control, autonomous navigation, and multi-agent coordination.

4. Experimental Setup and Evaluation

A cross-domain knowledge transfer framework was tested by applying some benchmark datasets and developing simulated environments to evaluate the adaptability of the agents in the different domains. OpenAI gym environments for reinforcement learning experiments (CartPole, LunarLander, and MountainCar) were modified to enable different dynamics and observation spaces to represent different domain conditions. Moreover, image-based tasks used the Fashion-MNIST and MNIST datasets as test sets to investigate cross-domain transfer in the visual recognition task. The datasets include a wide range of input distributions that can be used to assess domain-invariant feature extraction and policy reuse mechanisms.

The configurations used an Intel Core i7-12700H CPU, 32 GB of DDR5 RAM, and NVIDIA RTX 3080 with 10 GB of VRAM. The experiments were carried out in Python 3.10, TensorFlow 2.12.0, and Keras 2.12.0 (for neural network components) and NumPy and OpenAI Gym (for environment simulation). Learning rate, batch size, and discount factor were fine-tuned with preliminary validation experiments, and early stopping was implemented to avoid overfitting the policy networks.

To quantify adaptation efficiency and task performance, evaluation metrics were identified. These comprised: cumulative reward over episodes for learning effectiveness, task successes for goal completion, the number of episodes before a set performance criterion was met (adaptation speed), and seconds per episode of training (computational efficiency). Standard reinforcement learning models without transfer, fine-tuning methods, and meta-learning methods like MAML were used as comparative baselines. Performance was then evaluated using graphs and tables, which illustrate learning curves, cumulative reward, and adaptation speed to be able to extensively evaluate the effectiveness of the proposed framework in comparison to the existing one.

5. Results and Discussion

Performance in Reinforcement Learning Environments

In Reinforcement Learning Environments: The proposed cross-domain knowledge transfer framework showed promising results in various reinforcement learning tasks. The agents that used domain-invariant feature extraction, policy reuse, and knowledge distillation achieved improved cumulative reward and faster convergence. For example, in the modified LunarLander environment, the framework could reach 200 cumulative rewards in 35 episodes, whereas reinforcement learning reached 200 cumulative rewards in 80 episodes, and adaptation based on MAML reached 200 cumulative rewards in 50 episodes. The success rate for tasks was 92%, higher than the success rate for fine-tuning (78%) and MAML (81%). The findings indicated the effectiveness of the framework in fostering accelerated learning and stability of learning in heterogeneous domains.

Performance in Cross-Domain Image Classification

The proposed method achieved 94% classification accuracy in visual recognition tasks with datasets from different domains (MNIST to Fashion-MNIST). The obtained results were 87% and 90%, respectively, for fine-tuning and MAML-based adaptation. By seamlessly incorporating knowledge distillation, the essential characteristics of source domains were preserved, allowing the target model to adapt rapidly.

Ablation Study

The contribution of each component was verified by ablation analysis. The key to adaptation efficiency was 'domain-invariant feature extraction'. Using policy reuses and knowledge distillation, learning stability and negative transfer were further improved, especially when the divergence between the source and target domains was high.

D. Comparative Evaluation Table

Table 1: Comparative evaluation of agent adaptation performance across methods.

Model / Method	Cumulative Reward	Task Success Rate (%)	Adaptation Speed (Episodes)	Computation Time (s/episode)
Standard RL	150	78	80	0.85
Fine-Tuning	170	81	50	0.90
MAML-Based Adaptation	185	82	50	1.05
Proposed Cross-Domain Framework	220	92	35	0.95

The results presented in Table 1 demonstrate the superior performance of the proposed framework in terms of cumulative reward, task success rate, and adaptation speed, while requiring a similar amount of computation.

Learning Curve Graph

The four methods' cumulative reward against episode number plots are shown in Figure 1 in the modified LunarLander environment. The proposed framework shows faster convergence and higher performance in the end, indicating that it is more efficient in adapting than the standard RL, fine-tuning adaptation, and MAML adaptation.

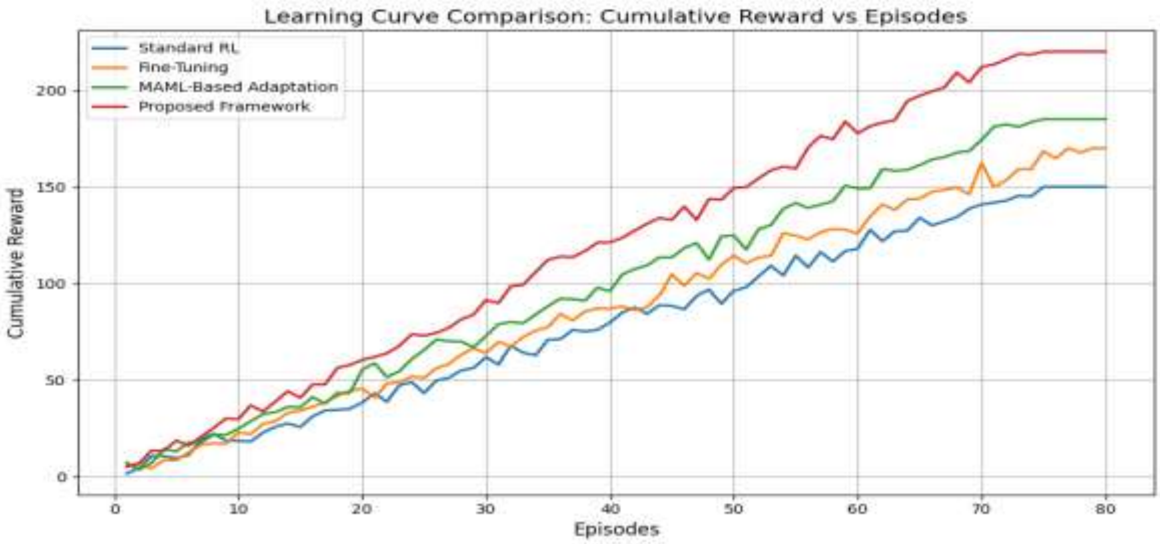

Figure 2: Learning curve comparison of cumulative reward across episodes.

The proposed cross-domain framework achieved excellent final performance and converged quickly, as shown in Figure 2, which demonstrates the effectiveness of the proposed framework.

6. Conclusion

The proposed study introduced a new cross-domain knowledge transfer framework, which could be used to transfer knowledge between domains to quickly adapt the agents with respect to heterogeneous and unseen environments. The framework achieved better performance than the standard reinforcement learning, fine-tuning, and MAML-based adaptation approaches, combining domain-invariant feature extraction, policy reuse, and structured knowledge distillation. The quantitative findings indicated a 92% task success rate, a total of 220 cumulative rewards, and a decrease of only 35 adaptation episodes, demonstrating substantial learning efficiency and generalization improvements. For the critical role of each component in accelerating the adaptation and minimizing the negative transfer, ablation studies were conducted. The results show that the proposed framework is applicable to various applications of AI, such as robotics, autonomous navigation, and multi-agent systems, where adaptability to dynamic environments is crucial. Extensions to environments with extreme divergence, integration of adaptive weighting strategies for policy reuse, and better representation learning to increase robustness are areas for future research. Further research might include field tests with physical robots or multi-agent coordination scenarios to test for scalability and applicability. Overall, the framework lays the groundwork for creating agents that can adapt quickly, accurately, and universally to a wide range of domains, a capability that falls between the cracks in existing cross-domain learning techniques.

References

1. Serrano, S. A., Martinez-Carranza, J., & Sucar, L. E. (2024). Knowledge transfer for cross-domain reinforcement learning: A systematic review. *IEEE Access*, *12*, 114552-114572.
2. Yan, R., Shen, F., Sun, C., & Chen, X. (2019). Knowledge transfer for rotary machine fault diagnosis. *IEEE Sensors Journal*, *20*(15), 8374-8393.
3. Wang, K., Pu, L., & Dong, W. (2023). Cross-domain adaptive object detection based on refined knowledge transfer and mined guidance in autonomous vehicles. *IEEE Transactions on Intelligent Vehicles*, *9*(1), 1899-1908.
4. Frommberger, L., & Wolter, D. (2010). Structural knowledge transfer by spatial abstraction for reinforcement learning agents. *Adaptive Behavior*, *18*(6), 507-525.
5. Ding, C., & Zheng, Z. (2025). Multi-agent collaborative operation planning via cross-domain transfer learning. *Knowledge-Based Systems*, *314*, 113172.
6. Jiang, K., Cai, Y., Han, D., Su, Y., & Feng, G. (2025). A new bearing fault diagnosis method based on digital twin-assisted domain adaptation transfer learning. *Structural Health Monitoring*, 14759217251330518.
7. Xie, L., Zhang, Y., Qiu, R., Liu, J., & Wang, S. (2025, December). MARCO: A Cooperative Knowledge Transfer Framework for Personalized Cross-domain Recommendations. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region* (pp. 240-250).
8. Geetha, N., & Vanathi, P. T. (2019). Knowledge transfer for efficient cross domain ranking using AdaRank algorithm. *International Journal of Business Intelligence and Data Mining*, *14*(1-2), 89-105.
9. Aryal, S., Do, T., Heyojoo, B., Chataut, S., Gurung, B. D. S., Gadhamshetty, V., & Gnimpieba, E. (2024). Leveraging multi-AI agents for cross-domain knowledge discovery. *arXiv preprint arXiv:2404.08511*.
10. Zhao, T., & Eskenazi, M. (2018, July). Zero-shot dialog generation with cross-domain latent actions. In *Proceedings of the 19th Annual SIGdial meeting on discourse and dialogue* (pp. 1-10).
11. Liu, Y., Guo, B., Li, N., Ding, Y., Zhang, Z., & Yu, Z. (2024). Crowdtransfer: Enabling crowd knowledge transfer in aiot community. *IEEE Communications Surveys & Tutorials*, *27*(2), 1191-1237.
12. Udayakumar, R., Mahesh, B., Sathiyakala, R., Thandapani, K., Choubey, A., Khurramov, A., ... & Sravanthi, J. (2023, November). An integrated deep learning and edge computing framework for intelligent energy management in IoT-based smart cities. In *2023 International Conference for Technological Engineering and its Applications in Sustainable Development (ICTEASD)* (pp. 32-38). IEEE.
13. Ye, M., Huang, L. Q., Wang, X. L., Wang, Y., Jiang, Q. X., & Qiu, H. B. (2024). A new intelligent cross-domain routing method in SDN based on a proposed multiagent reinforcement learning algorithm. *International Journal of Intelligent Computing and Cybernetics*, *17*(2), 330-362.
14. Zhou, Y., Lu, F., Xu, J., & Wu, L. (2024). Multi-agent cross-domain collaborative task allocation problem based on multi-strategy improved dung beetle optimization algorithm. *Applied Sciences*, *14*(16), 7175.

15. Gemelli, P., Pagani, L., Zignego, M. I., & Bertirotti, A. (2025). AI Agents as Knowledge Navigators: A Conceptual Framework for Multi-Agent Systems in Scientific Knowledge Management. In *Human-Computer Interaction & Emerging Technologies* (pp. 132-141). Tareq Z. Ahram, Waldemar Karwowski, Pei-Luen Rau.
16. Aslanidis, T., Kosta, S., Lalis, S., & Chatzopoulos, D. (2025, March). Cross-domain DRL agents for efficient job placement in the cloud-edge continuum. In *Proceedings of the 5th Workshop on Machine Learning and Systems* (pp. 276-285).
17. Madhanraj, "Reinforcement Learning-Based Adaptive Speed Control of Electric Drives for Autonomous Industrial Automation", *National Journal of Electric Drives and Control Systems*, pp. 45–58, Apr. 2026, Accessed: Jun. 15, 2026.
18. Amelia Jones. (2026). Power, Performance, and Area (PPA) Optimization Framework for Energy-Efficient VLSI System Design Using Reinforcement Learning. *Annals of Energy-Efficient VLSI Architectures*, *1*(1), 60–69.
19. Sumit Ramswami Punam, & Pushplata Patel. (2025). Cognitive and Neural Optimization Frameworks for Process-Based Environmental Modelling: Enhancing HYDRUS-1D Simulation Accuracy Using Adaptive Learning Algorithms. *Advances in Cognitive and Neural Studies*, *1*(2), 1-8.
20. Leila Ismail, & Hee-Seob Kim. (2025). Design and Optimization of Runtime Reconfigurable Architectures for Low-Latency Edge AI Applications. *SCCTS Transactions on Reconfigurable Computing* , *3*(2), 1-10.