

SHARING LIFELONG REINFORCEMENT LEARNING KNOWLEDGE VIA MODULATING MASKS

Saptarshi Nath¹, Christos Peridis¹, Eseoghene Ben-Iwhiwhu¹, Xinran Liu,²
Shirin Dora¹, Cong Liu³, Soheil Kolouri², Andrea Soltoggio¹

¹ Department of Computer Science, Loughborough University, United Kingdom

² Department of Electrical Engineering and Computer Science, Vanderbilt University, United States

³ Department of Computer Science and Engineering, University of California, Riverside, United States

{s.nath, c.peridis, e.ben-iwhiwhu, s.dora, a.soltoggio}@lboro.ac.uk,
{xinran.liu, soheil.kolouri}@vanderbilt.edu,
congl@ucr.edu

ABSTRACT

Lifelong learning agents aim to learn multiple tasks sequentially over a lifetime. This involves the ability to exploit previous knowledge when learning new tasks and to avoid forgetting. Recently, modulating masks, a specific type of parameter isolation approach, have shown promise in both supervised and reinforcement learning. While lifelong learning algorithms have been investigated mainly within a single-agent approach, a question remains on how multiple agents can share lifelong learning knowledge with each other. We show that the parameter isolation mechanism used by modulating masks is particularly suitable for exchanging knowledge among agents in a distributed and decentralized system of lifelong learners. The key idea is that isolating specific task knowledge to specific masks allows agents to transfer only specific knowledge on-demand, resulting in a robust and effective collective of agents. We assume fully distributed and asynchronous scenarios with dynamic agent numbers and connectivity. An on-demand communication protocol ensures agents query their peers for specific masks to be transferred and integrated into their policies when facing each task. Experiments indicate that on-demand mask communication is an effective way to implement distributed and decentralized lifelong reinforcement learning, and provides a lifelong learning benefit with respect to distributed RL baselines such as DD-PPO, IMPALA, and PPO+EWC. The system is particularly robust to connection drops and demonstrates rapid learning due to knowledge exchange.

1 INTRODUCTION

Distributed lifelong reinforcement learning (DLRL) is an emerging research field that offers scalability, robustness, and speed in real-world scenarios where multiple agents continuously learn. A key aspect of DLRL is the ability of multiple agents to learn multiple tasks sequentially while cooperating through information sharing. Notable advances have been made in related areas such as distributed reinforcement learning (Espeholt et al., 2018; Wijmans et al., 2019), federated reinforcement learning (Qi et al., 2021), and lifelong reinforcement learning (Zhan et al., 2017; Abel et al., 2018a;b; Xie et al., 2020). While the integration of lifelong learning and distributed learning is emerging only recently (Mohammadi & Kolouri, 2019; Song et al., 2023), the combination of these two paradigms could yield efficient, scalable, and robust learning systems.

The use of multiple workers in RL has been applied to accelerate data collection in methods like A3C (Mnih et al., 2016) while training a single model. Distributed reinforcement learning (DRL) (Weiß, 1995; Espeholt et al., 2018; Wijmans et al., 2019) expands on the concept of multiple workers to include multiple agents and may utilize a fully decentralized approach where each agent learns a potentially different model. Such systems can exhibit high learning performance and robustness. However, learning across multiple agents does not address the issue of non-IID data, which frequently arises when multiple tasks are learned sequentially over time.

Lifelong or continual learning specifically addresses the issue of sequential task learning, in which an agent learns multiple tasks characterized by unique input, transition, and reward distributions (Thrun, 1995; De Lange et al., 2021). The primary challenge is integrating a new task into existing knowledge without experiencing catastrophic forgetting

(McCloskey & Cohen, 1989). Various approaches have been suggested to implement lifelong learning: De Lange et al. (2021) propose a taxonomy with three main approaches, namely replay methods, regularization methods, and parameter isolation methods. Parameter isolation methods, in particular, allocate different parameters in a model to different tasks. In this paper, we suggest that such isolation mechanisms can be advantageous in a distributed setting where agents need to share knowledge.

We adopt a lifelong learning approach based on modulating masks, which have demonstrated competitive performance in supervised learning (Mallya & Lazebnik, 2018; Zhou et al., 2019; Wortsman et al., 2020; Koster et al., 2022). Recently, modulating masks have also been extended to lifelong reinforcement learning (LRL), enabling the integration and exploitation of previous knowledge when learning new tasks (Ben-Iwhiwhu et al., 2022b). A natural question arising from these recent studies is whether decomposing knowledge into masks can be utilized to transfer task-specific information across agents, implementing a distributed system of lifelong reinforcement learners. In this paper, we propose a system where agents use a fully asynchronous and decentralized protocol to query each other and exchange only relevant information for their respective tasks. The proposed proof-of-concept, named *lifelong learning distributed decentralized collective* (L2D2-C), suggests a lifelong learning approach to distributed decentralized learning. We test the proposed approach on benchmarks designed for multiple sequential tasks, namely the CT-graph (Soltoggio et al., 2023) and Minigrid (Chevalier-Boisvert et al., 2018). The simulations show that the system accelerates learning with respect to a single lifelong learning agent by a factor close to the number of agents. In addition, we demonstrate the system’s robustness to connection drops. Comparisons with existing distributed RL approaches, namely DD-PPO (Wijmans et al., 2019) and IMPALA (Espeholt et al., 2018), and non-distributed LL approaches (PPO+EWC) illustrate the advantages of integrating both LL and sharing. This study only illustrates initial experiments to suggest the validity of the idea, encouraging further future investigations for particular aspects of the system. The code to launch L2D2-C and reproduce the results is freely available at <https://github.com/DMIU-Shell/deeprl-shell>.

The paper is structured as follows. Section 2 provides a review of related work and introduces the necessary background concepts. Section 3 introduces the *lifelong learning distributed decentralized collective* (L2D2-C) approach. In Section 4, we present the results of our simulations and analyze them. Finally, in Section 5, we discuss the implications of our findings and conclude the paper. Additional information is provided in the Appendix.

2 RELATED WORKS AND BACKGROUND

We focus on a relatively unexplored scenario in which multiple lifelong learning agents learn from their own non-IID streams of tasks and share relevant acquired knowledge with each other upon request. While individual aspects of such a system have been studied in the literature, our work brings together several of these concepts to create a more comprehensive understanding. For example, lifelong reinforcement learning (Section 2.1) explores continual learning methods in RL settings for a single agent, while federated learning (Section 2.2) investigates the use of centralized and decentralized methods to learn with non-IID data. Finally, distributed RL (Section 2.3) seeks to enhance the speed of RL by deploying multiple agents working together. In addition, the concept of modulating masks has become crucial in various fields, including lifelong learning. In Section 2.4, we examine how modulating masks can be effectively employed to facilitate lifelong reinforcement learning.

2.1 LIFELONG REINFORCEMENT LEARNING

Lifelong learning, a concept known for decades (Thrun, 1995), has gained prominence as an active field of study in recent neural network studies (Soltoggio et al., 2018; Van de Ven & Tolias, 2019; Hadsell et al., 2020; Khetarpal et al., 2020; De Lange et al., 2021). The desiderata of a lifelong learner lie in its ability to learn multiple tasks in sequence while avoiding catastrophic forgetting (McCloskey & Cohen, 1989), enabling knowledge reuse via forward and backward transfer (Chaudhry et al., 2018) and reducing task interference (Kessler et al., 2022).

Different approaches can be largely categorized into regularization methods, modular methods, or a combination of both. Regularization methods apply an additional penalty term in the loss function to penalize either large structural changes or large functional changes in the network. Structural regularization or synaptic consolidation methods (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018; Kolouri et al., 2019) penalizes large change weights in the network that are important for maintaining the performance of previously learned tasks. Modular methods learn modules or a combination of modules useful for solving each task. A module is expressed as either: (i) a mask that activates sub-regions of a fixed neural network (Mallya & Lazebnik, 2018; Serra et al., 2018; Mallya et al., 2018; Wortsman et al., 2020; Koster et al., 2022; Ben-Iwhiwhu et al., 2022b), or (ii) a neural network that is expanded (Rusu et al., 2016; Yoon et al., 2018) or compositionally combined (Mendez & Eaton, 2021; Mendez et al., 2022) with other modules as new tasks are learned.

Lifelong RL agents are set up as a combination of standard deep RL and lifelong learning algorithms. In Kirkpatrick et al. (2017), EWC was combined with DQN (Mnih et al., 2013) to solve Atari games (Bellemare et al., 2013). Progress & Compress (Schwarz et al., 2018) showed further extensions of EWC by combining EWC with IMPALA (Espeholt et al., 2018) and policy distillation techniques. In CLEAR (Rolnick et al., 2019), IMPALA was combined with a replay method and behavioral cloning. SAC (Haarnoja et al., 2018) was combined with a number of lifelong learning algorithms in Wołczyk et al. (2021). For offline RL, Xie & Finn (2022) leveraged a replay buffer in SAC and importance weighting technique to specifically tackle forward transfer in a robotics agent, while Mendez et al. (2022) merged neural modular composition, offline RL, and PPO (Schulman et al., 2017) to enable reuse of knowledge and fast task learning. For masking approaches, Wołczyk et al. (2021) combined SAC and PackNet (Mallya & Lazebnik, 2018) (masks derived from iterative pruning), while Ben-Iwhiwhu et al. (2022b) demonstrated the use of directly learned modulating masks with PPO and introduced a knowledge reuse mechanism for the masking agents.

2.2 FEDERATED LEARNING FOR NON-IID TASKS

Federated learning (FL) generally considers multi-modal systems that train centralized models across various nodes, enabling accelerated training on massive amounts of data (Kairouz et al., 2021). Data anonymity and security are one of the core concerns in federated learning. FL approaches that address fully decentralized and non-IID scenarios are focused primarily on reducing the efficiency drops and guaranteeing convergence for a model for a single task. Sources of non-IID data in FL are often related to the geographical distributions of the clients, time of training, and other factors that relate often to biasing aspects of the data collection rather than the assumption of different tasks. As a result, rather than employing lifelong learning algorithms, more often FL on non-IID data adopts solutions to mitigate the effect of different distributions such as sharing sample data sets across clients (Wang et al., 2018; Zhao et al., 2018; Ma et al., 2022). An alternative FL approach to manage non-IID data is to assume the presence of multiple tasks: in such cases, FL integrates aspects of multi-task learning and meta-learning. In multi-task learning, the result of the process is one model per task (Smith et al., 2017; Zhang et al., 2022). Finally, meta-learning (Hospedales et al., 2021) can be effective to learn a global model for further fine-tuning to local data. Combinations of meta-learning approaches such as MAML (Finn et al., 2017) with FL have been explored in Jiang et al. (2019); Khodak et al. (2019). The objective of such approaches is often that of adapting a global model to a specific local data set (Fallah et al., 2020; Li et al., 2021; Tan et al., 2022) in what is referred to as personalized FL.

2.3 DISTRIBUTED REINFORCEMENT LEARNING

The concept of distributed reinforcement learning (DRL) was initially introduced to distribute computation across many nodes and thereby increase the speed at which data collection and training are performed (Weiß, 1995; Gronauer & Diepold, 2022). Both synchronous and asynchronous methods such as A2C/A3C (Mnih et al., 2016), DPPO (Heess et al., 2017), IMPALA (Espeholt et al., 2018), and others, use multiple workers to increase the rate of data collection while training a central model. A variation of PPO (Schulman et al., 2017), DD-PPO (Wijmans et al., 2019) extends the framework to be distributed and uses a decentralized optimizer, reporting a performance that scales nearly linearly with the number of nodes. While these algorithms have shown improvements in the SoTA on various benchmarks, they do not incorporate lifelong learning capabilities.

2.4 MODULATING MASKS FOR LIFELONG REINFORCEMENT LEARNING

The idea of using modulation in RL tasks is not new (Doya, 2002; Dayan & Niv, 2008; Soltoggio et al., 2008; Ben-Iwhiwhu et al., 2022a), but recently developed masking methods for deep supervised lifelong learning, e.g., Wortsman et al. (2020), have shown the advantages of isolation parameter methods. In Ben-Iwhiwhu et al. (2022b), modulating masks are shown to work effectively in deep reinforcement learning when combined with RL algorithms such as PPO (Schulman et al., 2017) or IMPALA (Espeholt et al., 2018). Other approaches such as Mallya & Lazebnik (2018) or Mallya et al. (2018) use masks in RL, but modify the backbone network as well, making them less suitable for knowledge sharing. In Ben-Iwhiwhu et al. (2022b) the agent contains a neural network policy $\pi_{\theta, \Phi}$, parameterized by the weights (backbone) of the network θ and the mask score parameters $\Phi = \{\phi_1, \dots, \phi_k\}$ for tasks $1 \dots k$. The backbone is randomly initialized (using the signed Kaiming constant (Ramanujan et al., 2020)) but kept fixed and remains unchanged, while mask score parameters are optimized during learning and applied on the backbone. For any given task k , ϕ_k represents the mask parameters across all layers of the network. In a network consisting of L layers, $\phi_k = \{S_k^1, \dots, S_k^L\}$ where S_k^i is the mask parameter in layer i for task k . To apply a mask on the backbone, ϕ_k is quantized using an element-wise threshold function (i.e., 1 if $\phi_{k, \{i, j\}} > 0$, otherwise 0) to generate a binary mask that is element-wise multiplied with θ . The multiplicative process activates and deactivates different regions of the backbone. Also, the framework supports knowledge reuse via the weighted linear combination of previously learned

task masks and the current task mask. For each layer l , this is expressed as:

$$S^{l,lc} = \left(\sum_{i=1}^k \beta_i^l S_i^{l,*} \right) + \beta_{k+1}^l S_{k+1}^l, \quad (1)$$

where $S^{l,lc}$ denotes the transformed mask score parameters for task $k + 1$ after the linear combination step. $S_i^{l,*}$ denotes the optimal mask score parameters for previously learned task i , and $\beta_1^l, \dots, \beta_{k+1}^l$ are the linear coefficients (weights) of the operation at layer l . Softmax is applied on the linear co-efficient to normalize it and also express the degree to which each mask is relevant to learn the current task. The learned parameters in the framework are β and Φ . Once the mask for task $k + 1$ has been learned using linear combination parameters, it is consolidated into a mask by itself, and thus the masks of other tasks can undergo further changes. The approach has been shown to be an effective way to implement LL by avoiding forgetting via complete parameter isolation and exploiting previous knowledge to facilitate learning of future tasks in a sequential curriculum.

3 SHARING MASKS ACROSS LIFELONG LEARNING AGENTS

The main research questions we want to address are, can we effectively transfer modulating masks across agents to produce a distributed lifelong learning system? Are masks a suitable form of parameter isolation that allows for task-specific policies to be transferred and integrated across agents? If we do so, can we observe a learning benefit from the synergy of both lifelong learning and sharing? The benefits may include increased learning speed of the collective with respect to the single agent, albeit at the increased overall computational resources. A hypothesis is that, if learning and sharing is efficient, n agents may be up to n -times faster at learning a given curriculum than the single agent. We provide further technical considerations on such an idea of performance improvement in Section A. To address the objectives above, we designed an algorithm for communication and sharing that can transfer the relevant modulating masks from the appropriate agents when needed. To produce a fully asynchronous and fully distributed system, each agent needs to act without central coordination and independently determine what information to ask for and when.

3.1 WHAT AND WHEN TO SHARE

We assume that each agent in the system shares the same backbone network θ that is randomly initialized. Each agent stores a list of masks $\Phi = \{\phi_1, \dots, \phi_k\}$ for tasks that they have encountered, and therefore they can independently learn a subset of all available tasks. A research question is whether agents can opportunistically and timely exchange learned masks to (i) acquire the knowledge of a task that has been already learned by another agent and integrate it into its own knowledge (ii) continue to learn on that task in cooperation with other agents, thereby increasing the knowledge of the collective system. With that aim, we introduce two main communication operations: queries and mask transfers.

Queries. When Agent 1 faces a task, it sends a query to all other agents (IDQ) communicating the task representation (e.g., task ID) and its performance on that task. Agents that have seen that particular task before, and exceed the performance of Agent 1, send a query-response (QR) with their performance measures, effectively communicating that they have better knowledge to solve that particular task. All other agents do not answer the query.

Mask transfers. Agent 1 then selects the best agent that reports the highest performance on the task and requests the mask with mask-request (MR) sent only to that particular agent, which, in turn, answers by sending the requested mask with a mask-transfer (MTR). In the present implementation, we transfer the scores of Eq. 1. Transferring the sparse and binarized masks offers significant bandwidth saving, but requires a re-initialization of the scores to resume training. Note that the scores are not used in forward passes, so such a process does not cause a drop in performance. The optimization of communication, including different degrees of sparsity of masks could be the topic of future application-specific studies.

In summary, when an IDQ is triggered, the following steps take place:

- Step 1: Agent 1 sends a query (IDQ) to all other agents with the task ID and its performance.
- Step 2: Agents that have encountered that task before and exceed the performance by at least 10% (estimated as the return over the last 512¹ steps), respond with a query-response (QR).
- Step 3: Agent 1 sorts all the responses to find the agent X that has the highest performance and requests a mask from that particular agent (via a mask-request (MR)).
- Step 4: Agent X sends the mask corresponding to that task ID back to agent 1 with a mask-transfer (MTR).

¹Assuming the product of the rollout length and number of workers in tables 4 and 5 is 512.

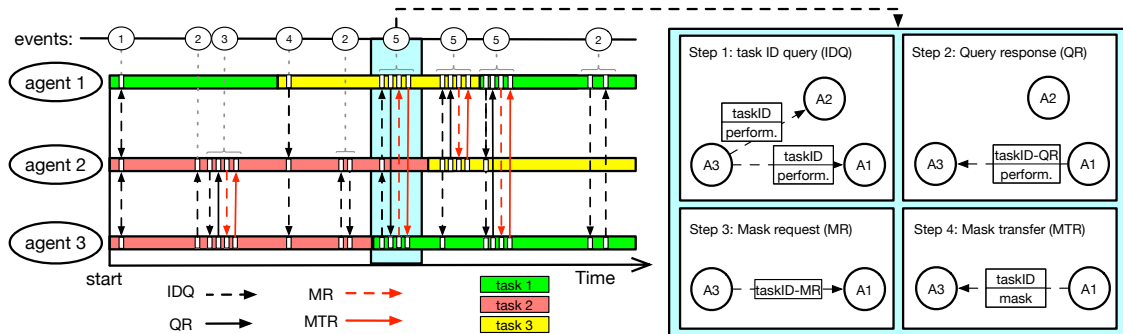


Figure 1: Messages and steps in the communication protocol. (Left) Event 1: All agents send IDQ to all other agents. Event 2: Agent 3 sends an IDQ to agent 2, there is no answer because agent 2 cannot exceed agent 3 performance. Event 3: An IDQ triggers a QR, followed by an MR and MTR. Event 4: Agent 1 starts a new task and issues IDQs, but there are no answers as no one has seen this task before. Event 5: Agent 3 starts task 1, issues IDQs, and receives a QR, which triggers an MR and MTR. (Right) A detailed description of event 4. An IDQ message is sent by agent 3 when starting to learn task 1. A QR is sent back by agent 1 who has seen task 1 before and exceeds agent 3 performance. Agent 3 then selects the best-performing policy and requests the corresponding mask to agent 1 with an MR message. In step 4, agent 1 sends the requested mask to agent 3 (MTR).

The process is graphically illustrated in Figure 1. An IDQ is triggered by one of two events: a task change or a maximum length of independent learning. In the first case (task change), an agent is about to start learning a task and therefore issues IDQs to consult the collective about whether a policy (mask) exists that solves the task better than its own mask. The second case, instead, is designed to take advantage of more frequent communication among agents learning the same task, which periodically check the collective for better policies than their own.

The number of IDQ messages sent across the network depends on the frequency f of IDQ messages per agent times n the number of agents, with a growth of fn^2 . The growth is quadratic with the number of agents in a fully connected network. However, such messages are small as they contain only the task ID. Moreover, they are sent only to reachable online agents. The number of QR messages is a fraction of the IDQ messages that depends on how many agents exceed the performance on that task. These are also small messages containing only the task ID and a measure of performance. The number of mask requests and responses grows linearly with the frequency of IDQs and the number of agents, i.e., fn . This is particularly important for the transfer of a mask, which is the largest type of communication that transfers the model parameters of Eq. 1.

3.2 KNOWLEDGE DISTRIBUTION, PERFORMANCE AND EVALUATING AGENTS

In a distributed system, determining how the collective’s knowledge is distributed across agents is not straightforward. The communication protocol devised in Section 3.1 implies that agents only maintain knowledge of the tasks that they have encountered, and update such knowledge from the collective only when they re-encounter that task. We devised the following system to assess the performance of the system as a whole. Given a curriculum $\mathcal{T} = \{\tau_1, \tau_2, \dots, \tau_n\}$, a measure of performance can be expressed as the degree to which an agent can solve those tasks after a given learning time μ . In other words, the learning objective can be set to maximize the reward across all tasks, i.e., a performance value p can be defined as:

$$p(\mu) = \sum_{\tau=1}^{\tau=n} \sum_{t=\mu}^{\mu+\kappa} r(s_t, a_t) \quad , \quad (2)$$

where μ is the time at which the system is assessed, κ is the duration of an evaluation block set as a multiple of an episode duration, and $r(s_t, a_t)$ is the reward collected at each time step. If we divide $p(\mu)$ by the number of episodes in an evaluation period κ (nr. of episodes in an evaluation ee), we obtain an instant cumulative return (ICR) of an agent at time μ , i.e., $ICR(\mu) = p(\mu)/ee$. Eq. 2, when computed over time also provides a basic but effective lifelong learning metric. While many other factors can be considered to assess lifelong learning systems (New et al., 2022; Baker et al., 2023), $p(\mu)$ in Eq. 2 reveals how well an agent can solve all tasks after a period of training.

To assess the performance of the learning collective with respect to a single agent, we also introduce two comparison values: (1) a performance advantage (PA) defined as the ratio $PA = p_c(\mu)/p_s(\mu)$, where p_c refers to a collective of agents and p_s refers to a single agent: this is the advantage of a collective of agents with respect to a single agent; (2)

a time advantage (TA), defined as the ratio between time-to-performance-p for the collective with respect to the single agent. This is how much faster the collective reaches a determined performance.

Eq. 2 requires the unrolling of the sequence of tasks \mathcal{T} to at least one agent in the collective. A common approach to test lifelong learning (Baker et al., 2023) is to stop learning to perform evaluation blocks, during which the agent is assessed. In our case, stopping any agent in the collective from learning in order to assess the performance will affect the performance itself. The solution we devised is to create one or more special agents called *evaluating agents* (EA) with the following characteristics: the EA does not learn and does not answer queries, so it is invisible to the collective and it cannot impact their performance. However, it can query all agents and fetch their knowledge, which is then tested continuously. The use of EAs is a way to monitor the performance of the collective without interfering with its learning dynamics. In the experiments in which we used the ICR metric (derived by Eq. 2), one evaluation agent was deployed.

3.3 DECENTRALIZED DYNAMIC NETWORK AND AGENT ARCHITECTURE

To ensure that L2D2-C is completely decentralized and asynchronous, the entire architecture is contained within one agent, i.e., there is no coordinating or central unit. Each agent in the system is identified by their unique tuple $\langle \text{IP}, \text{Port} \rangle$. An agent stores the following data structures: *entry-points* and *online-agents*. *Entry-points* is a list of $\langle \text{IP}, \text{Port} \rangle$ tuples that lists all known agents. The list is continuously updated by adding any new agent that makes contact via an IDQ. When an agent enters the collective and makes first contact with another agent, it receives the *entry-points* list from that agent, thus acquiring knowledge of the identities of the other agents. If an agent goes offline for a period of time, when is back online, it will attempt to access the collective by contacting the agents in this list. The *online-agents* list keeps track of agents that are currently online, and it is used to send out IDQs. This system allows agents to (i) dynamically enter and leave the collective at any time and (ii) manage sparsely connected or constrained topologies by communicating only with reachable (i.e. online) agents. Such a simple setup allows L2D2-C to run multiple agents on the same server, across multiple servers, and across multiple locations anywhere in the world.

As L2D2-C emerges from instantiating multiple single agents, the entire system architecture is contained within one agent. In the most general case of DLRL, events such as task changes, data collection and learning, and communication among agents, are stochastic and independent events. Thus, an agent is required to perform different operations in parallel (data collection, learning, and communication as both client and server) to maximize performance. L2D2-C is implemented with multi-processing and multi-threading to ensure all such operations can run simultaneously. An overview of the agent’s architecture is provided in Figure 6 in the Appendix B.

4 EXPERIMENTS

We present the tests of L2D2-C when multiple agents learn randomized curricula of tasks. The following metrics are measured: performance of a varying number of agents with respect to the single lifelong learning agent (Section 4.2); lifelong learning performance of L2D2-C with respect to IMPALA, DD-PPO, PPO, PPO+EWC (single-head) and PPO+EWC (multi-head) (Section 4.3); performance of L2D2-C with unreliable communication (Section 4.4).

4.1 BENCHMARKS

We employed two benchmarks for discrete RL scenarios that are particularly suitable for sequential LL, the configurable tree graph (CT-graph) (Soltoggio et al., 2023) (code at Soltoggio et al. (2019)) and the Minigrid (Chevalier-Boisvert et al., 2018) environments. The CT-graph is a generative environment that implements a tree graph with 2D images which makes it suitable for evaluating lifelong RL approaches. A number of tasks with varying complexity can be automatically generated to create a large variety of curricula. It features sparse rewards and arbitrarily long episodes in different configurations. Due to the varying reward location, the CT-graph implements interfering tasks. Graphical illustrations of the environment are provided in Appendix D.1. The Minigrid (Chevalier-Boisvert et al., 2018) is a grid-world navigation environment, consisting of a number of predefined partially observable tasks with varying levels of complexity. The agent is required to reach a defined location in the grid world, avoiding obstacles such as walls, lava, moving balls, etc. The experiment protocol employs a curriculum of three tasks that consists of the following: SimpleCrossingS9N1, SimpleCrossingS9N2, SimpleCrossingS9N3. Screenshots of all tasks are reported in the Appendix (D.2).

4.2 L2D2-C VERSUS A SINGLE LIFELONG LEARNER (L2D2-C WITH ONE AGENT)

Our initial tests aim to assess the impact on performance with an increase in the number of agents in a LL setting. In the CT-graph, 16 tasks are learned sequentially over 64 learning slots per agent. One learning slot corresponds to

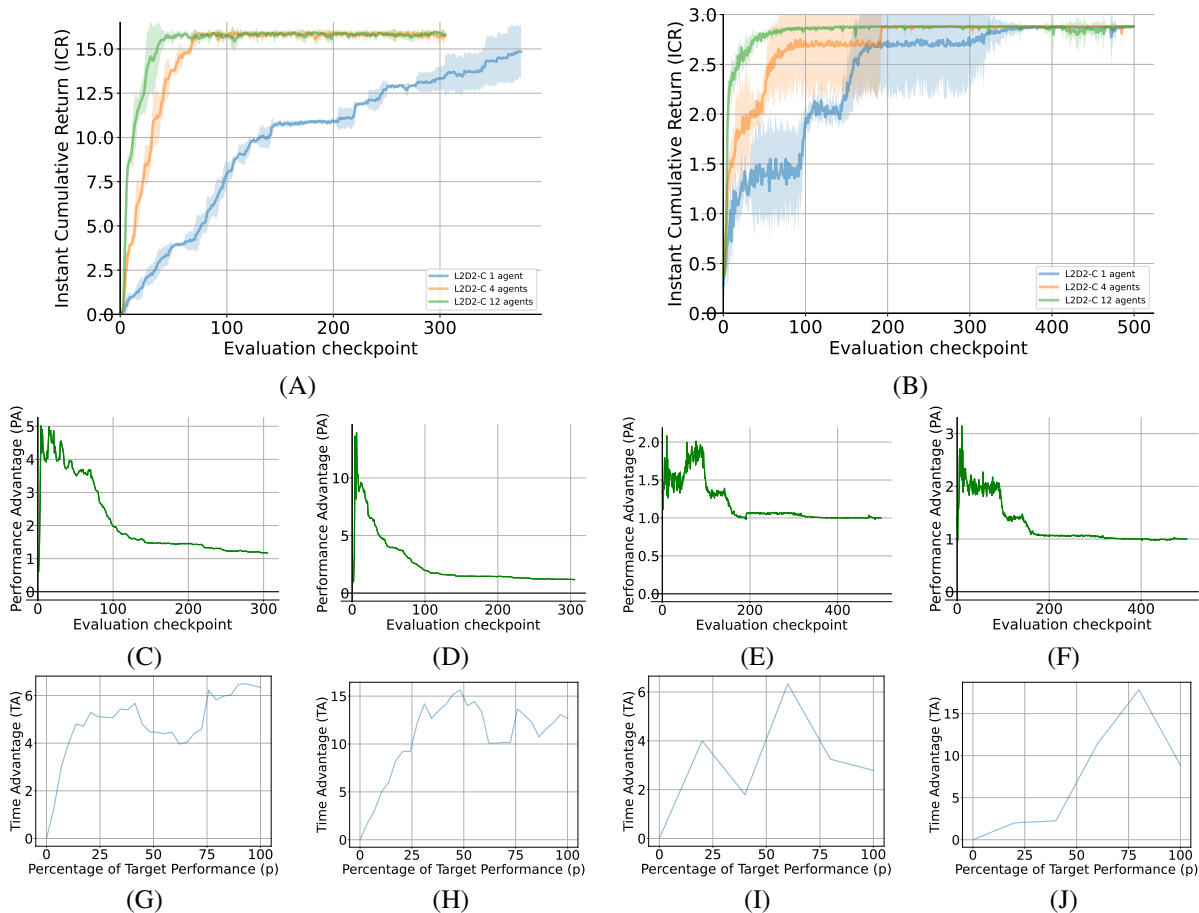


Figure 2: L2D2-C versus a single agent. (A) The instant cumulative return (ICR) provides the sum of the performance of the evaluating agents on 16 CT-graph tasks. A single agent is compared with a 4 and a 12-agent collective. (B) ICR on 3 Minigrid tasks. (C, D, E, F) Performance advantage (PA) of L2D2-C with respect to the single agent: (C) 4-agent, CT-graph; (D) 12-agent, CT-graph; (E) 4-agent, Minigrid; (F) 12-agent, Minigrid. (G, H, I, J) The time advantage (TA), i.e., the speedup of L2D2-C with respect to the single agent, expressed as a ratio of time-to-performance: (G) 4-agent, CT-graph; (H) 12-agent, CT-graph; (I) 4-agent, Minigrid; (J) 12-agent, Minigrid.

continuous training over one task for 12800 steps. In the Minigrid, 3 tasks are learned over 12 learning slots of 102400 steps each. The instant cumulative return (ICR) (Section 3.2) is used as the performance metric that represents the sum of returns across all tasks, i.e., $p(\mu)/ne$ (Eq. 2). The results in Figure 2 are the mean of 5 seed runs with the shades denoting the 95% confidence interval, computed following the procedure in Colas et al. (2018).

As the learning takes place, one evaluation agent (EA) was deployed to run a continuous assessment of L2D2-C by continuously cycling through the 16 tasks (CT-graph) and 3 tasks (Minigrid). As each agent runs independently from the others, including the EA, there is no direct mapping between evaluation checkpoints and training steps, as each agent could run at a slightly different speed. However, as the EA runs at a constant speed, the evaluation blocks can be interpreted as elapsed time. The graph indicates that the main advantage of L2D2-C is visible with a higher number of tasks (16 in the CT-graph), while such an advantage is reduced when learning the 3-task Minigrid. Nevertheless, the 12-agent L2D2-C appears immune from the local minima that affect the single agent and the 4-agent L2D2-C (Fig. 2(B)). The second and third row of Fig. 2 shows how much more performance is scored by L2D2-C with respect to the single agent (Panels C to F) and how much faster is L2D2-C to reach a determined level of performance (Panels G to J). From the second row of Fig. 2, the 4 and 12-agent L2D2-C have a significant performance advantage during the early stages of training: for the 16-task experiment, such an advantage peaks at levels that are comparable to the number of agents, i.e., approximately n times performance (where n is the number of agents), but decreases progressively with time as the single agent, slowly, learns all available tasks. The third row of Fig. 2 shows that the speed advantage is




Figure 3: Training performance on single tasks through a LL curriculum: two L2D2-C agents vs. two DD-PPO workers, two IMPALA workers, one PPO worker, one PPO+EWC single-head and one PPO+EWC multi-head worker. The performance of one single agent is plotted for each algorithm, averaged over 5 seeds and with the shade showing the 95% confidence interval.

also significant in the 16-task experiment and mostly exceeding a factor n for target performances above $\sim 30\%$, but not so for the 3-task experiment in which the low number of tasks per agent appears to have an impact on efficiency.

In addition to the experiments in Fig. 2, we run tests with the following agents/tasks numbers: 2/2, 4/4, 8/8, 4/16, 16/16, and 32/32. As these runs were executed with one seed, we report the results in the Appendix C in Table 9 and Fig. 8 for the 32-tasks/32-agents run as supporting experiments. The additional runs confirm the learning dynamics of Fig. 2 and indicate similar learning trends with up to 32 agents and 32 tasks.

4.3 L2D2-C VERSUS DD-PPO, IMPALA, PPO, PPO+EWC

We assessed the advantage of introducing LL dynamics to distributed RL, and in particular the performance of L2D2-C with respect to established distributed RL baselines, DD-PPO, IMPALA, plus two additional references PPO and PPO+EWC. Approaches were tested on a randomized 4-task CT-Graph curriculum with 3 repeats totaling 12 tasks at 42752 steps per task (Details in the appendix table 4). Rather than measuring the ICR metric that is implemented only for L2D2-C, we monitor the performance on each task individually for both L2D2-C and the baselines. Fig. 3 shows L2D2-C and the baselines sequentially learning different tasks. The task changes are visible when there is a drop in performance. It can be noted that an L2D2-C agent learns three tasks initially and then shows minor drops in performance as (1) it acquires task knowledge from the other agent and (2) it shows no forgetting thanks to the LRL masking approach. The non-LL baselines, instead, while showing steep learning curves, are unable to maintain the knowledge when switching tasks on such a sequential LL curriculum. The LL baseline PPO+EWC multi-head shows the best performance among other approaches, but such an algorithm does not have a sharing mechanism for distributed learning.

4.4 ROBUSTNESS IN FRONT OF CONNECTION DROPS

We introduced a stochastic mechanism to simulate connection drops: with varying probability, Step 1 in the communication protocol (issuing of IDQ as outlined in Section 3.1) is canceled. Since IDQ issuing is the first step in the communication chain, canceling this first step has the same effect as dropping communication at any following step. We use the evaluation agent in two different settings: to monitor all agents, and to monitor one single random agent. The evaluation agent is not affected by connection drops, it cannot learn, it cannot share and its presence does not affect the performance of L2D2-C. Fig. 4 shows that even high levels of connection drops fail to affect the performance of the system significantly until communication is completely stopped.

Where does such robustness come from? We identified two reasons. Firstly, when a connection is dropped, the agent continues to learn and attempts to communicate again sometime later (when the conditions for communication occur again). In other words, the agent gives up temporarily, uses the time to learn, and tries again later. Thus, we measured that the number of transferred masks (Table 1) does not decrease linearly with the probability of connection drops, and goes to zero only when communication is completely stopped. Detailed scatter plots of data exchange are reported in the appendix (Sec. C, Fig. 10). A second reason for robustness is that, even when communication is very unlikely or not possible at all, each agent is unhindered in their lifelong learning dynamics. Thus, many agents learning different tasks will collectively acquire more knowledge, which can then be shared the moment communication is restored. To

Type of message	Average nr. sent (5 seeds)				
	0% drop	50% drop	75% drop	95% drop	100% drop
ID queries (IDQ)	2705.98	1356.217	674.47	136.45	0
Query responses (QR)	187.23	95.25	51.33	18.82	0
Mask requests (MR)	28.40	27.45	25.03	15.88	0
Mask transfers (MTR)	28.40	27.45	25.03	15.88	0

Table 1: Average number of messages sent per type per agent with different probabilities of connection drops. It can be noted that even with high probabilities of connection drops, the agents “insist” on their attempts so that the number of masks transferred decreases to a lesser degree. The data shown above corresponds to the duration of learning in the collectives depicted between evaluation checkpoints 0 to 154 Fig. 4 (A) and 0 to 192 Fig. 4 (B).

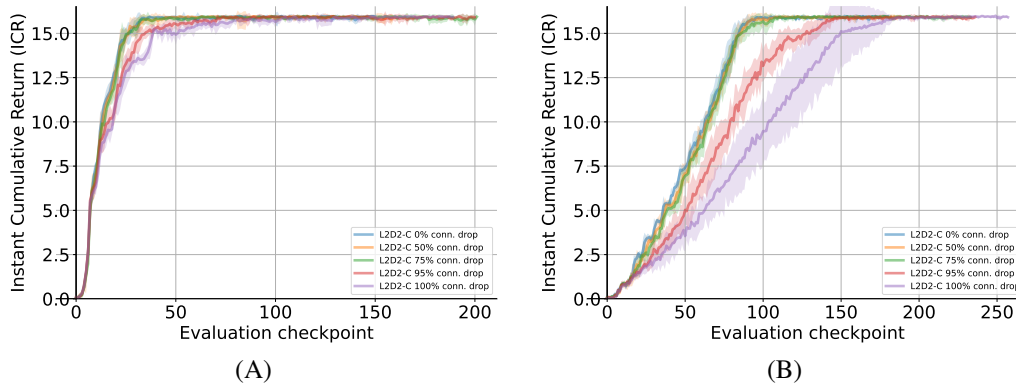


Figure 4: Instant cumulative reward (ICR) of L2D2-C with different connection drop probabilities for a 16-task, 12-agent CT-graph experiment (average and 95% confidence interval over 5 seeds). (A) One evaluation agent monitors all agents. (B) One evaluation agent monitors one random agent only.

provide a visual insight into learning for a single agent, Fig. 5 shows how a single agent learns a sequence of tasks when has full communication, 50%, and 100% connection drops.

5 DISCUSSION

In our experiments, we assumed that each agent sees a random task from a list of tasks and will learn that for a fixed duration. Under such conditions, some agents might be learning the same task at the same time, while other agents learn different tasks. The ratio of tasks/agents determines the similarity of DLRL with other approaches. For one-task/many-agents, LL is not necessary and DRL performs optimally. For many-tasks/one-agent, only LL is required. For many-tasks/many-agents a combination of DRL and LL is required, as in the proposed approach. In the experiments of Fig. 2, we investigated two cases. In the first case (CT-graph), we used more tasks than agents (16 tasks with 12, 4, and 1 agents) and saw a linear improvement of the L2D2-C metrics with the number of agents (Fig. 2 (CDGH)). In the second case (Minigrid), we used fewer tasks than agents (3 tasks with 12, 4, and 1 agents) and noticed the ability of parallel search to escape local minima as the number of agents increases (Fig. 2 (B)).

L2D2-C requires no central coordination with the exception of the assignment of the task ID, for which we assume the existence of an oracle to provide each agent with the ID for each task. Such an assumption could be unrealistic in real-world scenarios where tasks are not defined by an oracle (Rios & Itti, 2020). Extensions of the L2D2-C framework could be considered by adding task detection capabilities to enable agents to compare tasks according to their similarities (Liu et al., 2022). Such added capabilities would affect the queries that agents exchange (steps 1 and 2 in Section 3.1) that would therefore contain task descriptors as opposed to task IDs.

Further analysis of the system is required to assess the overall communication load and the impact of different network topologies. The query messages (IDQ and QR, Section 3.1) are designed to discover which agents have knowledge of a specific task, and therefore, their number grows with the square of the number of agents in a fully connected network. For larger networks, query messages could be optimized by adopting multiple stellar topologies, i.e., defining some agents as hubs. However, once the mask provider has been identified, our one-to-one mask transfer protocol ensures scalability. The impact of network topologies and the communication load was not presented in this paper as the experiments in this first study were run on a single server. However, the agent-centered L2D2-C architecture was




Figure 5: Examples of training plots for one agent with different levels of connection drops. Each task is learned for 10 iterations, equivalent to 12,800 steps. To improve readability, we modified the curriculum for this specific agent to undertake tasks sequentially from 1 to 16, and repeat that sequence. (A) Full communication (no drops). The upwards trend from about half performance shows that the agent receives masks and then improves on them. (B) 50% connection drops. (C) No communication. The masks now are trained partially when tasks are seen for the first time (iterations 1 to 160). The partially trained masks are used to continue training after iteration 160. No communication means that this agent is learning “on its own”.

observed to be functional across distant locations with a proof-of-concept experiment in which four agents formed an L2D2-C system across locations in two different countries. Under such conditions, analysis of the effect of bandwidth usage and delay is essential to further assess the system.

The current approach relies on a common backbone network shared initially by all agents, and therefore it cannot be used if agents have different networks. Other approaches, e.g., based on knowledge distillation (Hinton et al., 2015; Gou et al., 2021), could be devised for heterogeneous agents.

Important considerations for specific applications of the system could include computational and memory costs. As in all deep network approaches, these directly relate to the size of the network. In addition, by transferring highly sparse binary masks, it is possible to significantly reduce the ratio of the size of the mask with respect to the backbone. For example, for a backbone encoded with 32-bit precision parameters, a full binary mask is 1/32 the size of the backbone. Specific applications might result in different trade-offs of performance versus computation and memory. We did not implement optimization steps in the current study.

6 CONCLUSION

This paper introduced the concept of a distributed and decentralized RL system that can learn multiple tasks sequentially without forgetting thanks to lifelong learning dynamics. The system is based on the idea of agents exchanging task-specific modulating masks that are applied to a backbone network that is common to all the agents in the system. While the advantage of modulating masks is known in the literature for both supervised and reinforcement learning approaches, here we show that the isolation of task knowledge to masks can be exploited to implement a fully distributed and decentralized system that is defined simply by the interconnection of a number of identical agents. The L2D2-C system was shown to maintain LL dynamics across multiple agents and to prevent catastrophic forgetting that is typical of distributed approaches such as DD-PPO and IMPALA. The system has an increase in speed-up and performance that appears to grow linearly with the number of agents when the agent/task ratio is close to one. Finally, we observed a surprising robustness of the collective learning in front of high levels of connection drops.

BROADER IMPACT

The idea that reinforcement learning agents can learn sequential tasks incrementally and in collaboration with other agents contributes towards the creation of potentially more effective and ubiquitous RL systems in a variety of application scenarios such as industrial robotics, search and rescue operations, and cyber-security just to mention a few. A broad application of such systems could affect efficiency and productivity by introducing automation and independent ML-driven decision-making. Careful considerations must be taken in all scenarios in which human supervision is reduced to ensure the system acts in alignment with regulations, safety protocols, and ethical principles.

ACKNOWLEDGMENTS

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract Contract No. HR00112190132 (Shared Experience Lifelong Learning).

REFERENCES

- David Abel, Dilip Arumugam, Lucas Lehnert, and Michael Littman. State abstractions for lifelong reinforcement learning. In *International Conference on Machine Learning*, pp. 10–19. PMLR, 2018a.
- David Abel, Yuu Jinnai, Sophie Yue Guo, George Konidaris, and Michael Littman. Policy and value transfer in lifelong reinforcement learning. In *International Conference on Machine Learning*, pp. 20–29. PMLR, 2018b.
- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 139–154, 2018.
- Anyscale-Inc. Ray. <https://docs.ray.io/en/master/index.html>, May 2023.
- Megan M. Baker, Alexander New, Mario Aguilar-Simon, Ziad Al-Halah, Sébastien M.R. Arnold, Ese Ben-Iwhiwhu, Andrew P. Brna, Ethan Brooks, Ryan C. Brown, Zachary Daniels, Anurag Daram, Fabien Delattre, Ryan Del-lana, Eric Eaton, Haotian Fu, Kristen Grauman, Jesse Hostetler, Shariq Iqbal, David Kent, Nicholas Ketz, Soheil Kolouri, George Konidaris, Dhireesha Kudithipudi, Erik Learned-Miller, Seungwon Lee, Michael L. Littman, Sandeep Madireddy, Jorge A. Mendez, Eric Q. Nguyen, Christine Piatko, Praveen K. Pilly, Aswin Raghavan, Abrar Rahman, Santhosh Kumar Ramakrishnan, Neale Ratzlaff, Andrea Soltoggio, Peter Stone, Indranil Sur, Zhipeng Tang, Saket Tiwari, Kyle Vedder, Felix Wang, Zifan Xu, Angel Yanguas-Gil, Harel Yedidsion, Shangqun Yu, and Gautam K. Vallabha. A domain-agnostic approach for characterization of lifelong learning systems. *Neural Networks*, 2023. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2023.01.007>. URL <https://www.sciencedirect.com/science/article/pii/S0893608023000072>.
- Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Eseoghene Ben-Iwhiwhu, Jeffery Dick, Nicholas A Ketz, Praveen K Pilly, and Andrea Soltoggio. Context meta-reinforcement learning via neuromodulation. *Neural Networks*, 152:70–79, 2022a.
- Eseoghene Ben-Iwhiwhu, Saptarshi Nath, Praveen K Pilly, Soheil Kolouri, and Andrea Soltoggio. Lifelong reinforcement learning with modulating masks. *arXiv preprint arXiv:2212.11110*, 2022b.
- Arslan Chaudhry, Puneet K Dokania, Thalaisyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 532–547, 2018.
- Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. How many random seeds? statistical power analysis in deep reinforcement learning experiments. *arXiv preprint arXiv:1806.08295*, 2018.
- Peter Dayan and Yael Niv. Reinforcement learning: the good, the bad and the ugly. *Current opinion in neurobiology*, 18(2):185–196, 2008.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.
- Kenji Doya. Metalearning and neuromodulation. *Neural networks*, 15(4-6):495–506, 2002.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International conference on machine learning*, pp. 1407–1416. PMLR, 2018.
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948*, 2020.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129:1789–1819, 2021.

- Sven Gronauer and Klaus Diepold. Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, pp. 1–49, 2022.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.
- Nicolas Heess, Dhruva TB, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021.
- Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019.
- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning, 2021.
- Samuel Kessler, Jack Parker-Holder, Philip Ball, Stefan Zohren, and Stephen J. Roberts. Same state, different task: Continual reinforcement learning without interference, 2022.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *arXiv preprint arXiv:2012.13490*, 2020.
- Mikhail Khodak, Maria-Florina F Balcan, and Ameet S Talwalkar. Adaptive gradient-based meta-learning methods. *Advances in Neural Information Processing Systems*, 32, 2019.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Soheil Kolouri, Nicholas A Ketz, Andrea Soltoggio, and Praveen K Pilly. Sliced cramer synaptic consolidation for preserving deeply learned representations. In *International Conference on Learning Representations*, 2019.
- Nils Koster, Oliver Grothe, and Achim Rettinger. Signing the supermask: Keep, hide, invert. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=e0jtGTfPihs>.
- Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pp. 6357–6368. PMLR, 2021.
- Xinran Liu, Yikun Bai, Yuzhe Lu, Andrea Soltoggio, and Soheil Kolouri. Wasserstein task embedding for measuring task similarities. *arXiv preprint arXiv:2208.11726*, 2022.
- Zichen Ma, Yu Lu, Wenye Li, and Shuguang Cui. Efl: Elastic federated learning on non-iid data. In Sarath Chandar, Razvan Pascanu, and Doina Precup (eds.), *Proceedings of The 1st Conference on Lifelong Learning Agents*, volume 199 of *Proceedings of Machine Learning Research*, pp. 92–115. PMLR, 22–24 Aug 2022. URL <https://proceedings.mlr.press/v199/ma22a.html>.

- Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2018.
- Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 67–82, 2018.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pp. 109–165. Elsevier, 1989.
- Jorge A Mendez and Eric Eaton. Lifelong learning of compositional structures. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=ADWd4TJO13G>.
- Jorge A Mendez, Harm van Seijen, and Eric Eaton. Modular lifelong reinforcement learning via neural composition. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=5XmLzds1FNN>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937. PMLR, 2016.
- Javad Mohammadi and Soheil Kolouri. Collaborative learning through shared collective knowledge and local expertise. In *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6. IEEE, 2019.
- Alexander New, Megan Baker, Eric Nguyen, and Gautam Vallabha. Lifelong learning metrics. *arXiv preprint arXiv:2201.08278*, 2022.
- Jiaju Qi, Qihao Zhou, Lei Lei, and Kan Zheng. Federated reinforcement learning: techniques, applications, and open challenges. *Intelligence & Robotics*, 2021. doi: 10.20517/ir.2021.02. URL <https://doi.org/10.20517%2Ffir.2021.02>.
- Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11893–11902, 2020.
- Amanda Rios and Laurent Itti. Lifelong learning without a task oracle. In *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 255–263. IEEE, 2020.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In *International Conference on Machine Learning*, pp. 4528–4537. PMLR, 2018.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pp. 4548–4557. PMLR, 2018.
- Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/6211080fa89981f66b1a0c9d55c61d0f-Paper.pdf>.

- Andrea Soltoggio, John A Bullinaria, Claudio Mattiussi, Peter Dürri, and Dario Floreano. Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios. In *Proceedings of the 11th international conference on artificial life (Alife XI)*, number CONF, pp. 569–576. MIT Press, 2008.
- Andrea Soltoggio, Kenneth O Stanley, and Sebastian Risi. Born to learn: the inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks*, 108:48–67, 2018.
- Andrea Soltoggio, Pawel Ladosz, Eseoghene Ben-Iwhiwhu, and Jeff Dick. The CT-graph environments, 2019. URL <https://github.com/soltoggio/ct-graph>.
- Andrea Soltoggio, Eseoghene Ben-Iwhiwhu, Christos Peridis, Pawel Ladosz, Jeffery Dick, Praveen K Pilly, and Soheil Kolouri. The configurable tree graph (ct-graph): measurable problems in partially observable and distal reward environments for lifelong reinforcement learning. *arXiv preprint arXiv:2302.10887*, 2023.
- Ziang Song, Zhuolong Yu, Jingfeng Wu, Lin Yang, and Vladimir Braverman. Poster: Fledge: Federated lifelong learning on edge devices. <https://crossf12022.github.io/abstracts/Abstract5.pdf>, 2023. Accessed: 2023-03-13.
- Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- Sebastian Thrun. A lifelong learning perspective for mobile robot control. In *Intelligent robots and systems*, pp. 201–214. Elsevier, 1995.
- Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.
- Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.
- Gerhard Weiß. Distributed reinforcement learning. In *The Biology and technology of intelligent autonomous agents*, pp. 415–428. Springer, 1995.
- Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv preprint arXiv:1911.00357*, 2019.
- Maciej Wołczyk, Michał Zając, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Continual world: A robotic benchmark for continual reinforcement learning. *Advances in Neural Information Processing Systems*, 34:28496–28510, 2021.
- Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *Advances in Neural Information Processing Systems*, 33:15173–15184, 2020.
- Annie Xie and Chelsea Finn. Lifelong robotic reinforcement learning by retaining experiences. In *Conference on Lifelong Learning Agents*, pp. 838–855. PMLR, 2022.
- Annie Xie, James Harrison, and Chelsea Finn. Deep reinforcement learning amidst lifelong non-stationarity. *arXiv preprint arXiv:2006.10701*, 2020.
- Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Sk7Ksfw0->.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pp. 3987–3995. PMLR, 2017.
- Yusen Zhan, Haitham Bou Ammar, and Matthew E Taylor. Scalable lifelong reinforcement learning. *Pattern Recognition*, 72:407–418, 2017.
- Hongwei Zhang, Meixia Tao, Yuanming Shi, and Xiaoyan Bi. Federated multi-task learning with non-stationary heterogeneous data. In *ICC 2022-IEEE International Conference on Communications*, pp. 4950–4955. IEEE, 2022.
- Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. *Advances in neural information processing systems*, 32, 2019.

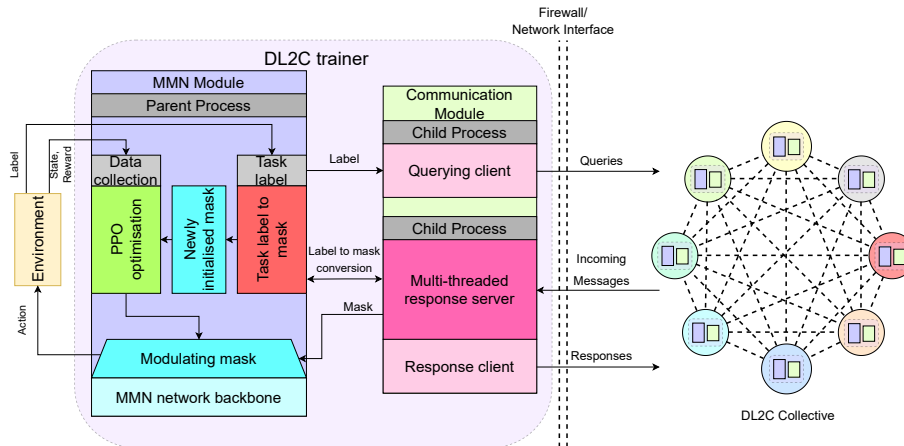


Figure 6: System architecture of an L2D2-C agent. Collectives consist of multiple instances with a common network backbone.

A CONSIDERATION ON THE PERFORMANCE OF n AGENTS

With respect to a single-agent system, in a multi-agent system, the overall computational cost increases by a factor that is proportional to the number of agents, plus a communication cost. However, the aim of deploying multiple agents is to increase the speed of learning and observe possible beneficial effects of inter-agent interaction. When comparing the performance of n agents versus one, it might appear as “unfair”, but the point of the comparison is to observe how much faster can we learn a set of tasks if n agents are deployed instead of one. In fact, it is not given that n agents will learn n times faster than the single agent: if n agents do not communicate, each agent is effectively a single agent and only the merging of their knowledge would cause an acceleration of learning. If we define a metric “learning speed (LS)” as the time taken by an agent to reach a pre-determined level of performance on a given curriculum of tasks, we can distinguish the following cases:

- $LS(n) \sim LS(1)$: n agents take the same amount of time to learn the curriculum as one agent. This is the case when agents do not communicate and their knowledge is not aggregated. The cost of n agents is not exploited.
- $1/n \cdot LS(1) < LS(n) < LS(1)$: n agents learn faster than one agent but are less than n times faster than the single agent. This is the case when agents benefit from communication, but there is a loss in efficiency. The cost of n agents is somewhat compensated by improved performance.
- $LS(n) \sim LS(1)/n$: n agents are approximately n times faster than the single agent. In this case, the increased computational cost of a factor n is compensated by an equivalent improvement in performance.
- $LS(n) < LS(1)/n$: n agents are faster than n times the single agent. While this may appear unlikely at first, such a situation can emerge if n agents exploit parallel search to solve difficult problems, e.g., with sparse reward. In this case, the cost of n agents is largely compensated by a better-than- n improvement in performance.

Overall, we can conclude that the deployment of n agents versus one may lead to significant gains in performance if an efficient algorithm can exploit their combined computation.

B HARDWARE AND SOFTWARE DETAILS

Simulations were conducted with two main network setups: one with all agents launched on a single GPU server, with specifications described in Table 2 and a second setup in which agents were launched on different servers in different locations.

Fig. 7 illustrates the steps of the communication protocol (Algorithms 1 and 2) in two flowcharts.

The code to run L2D2-C and reproduce the experiment presented in this paper is available at <https://github.com/DMIU-ShELL/deeprl-shell>.




Figure 7: Flowcharts depicting the communication protocol used in L2D2-C. Left: Client-side, Right: Server-side

Algorithm 1 Client Protocol**Require:** List of IP-port tuples, C **Require:** Current task label, L **Require:** Current task reward, R **Require:** Constant identifier for query, I **Require:** Own IP address, ip_{self} **Require:** Own port, $port_{self}$

```

1: while True do
2:   while Label  $L$  not received from the main process do
3:     Wait for label  $L$ 
4:   end while
5:   Construct query as array  $\{ip_{self}, port_{self}, I, L, R\}$ 
6:   Serialize query into a byte string
7:   for  $(ip, port)$  in  $C$  do
8:     Send outgoing query
9:   end for
10:  Sleep for 0.2 seconds
11:  if Length of response list  $> 0$  then
12:    Select best agent  $A_n$  using condition  $A_n = \{a \in C : 0.9R_n > R_x\}$  from received responses
13:    Send mask request to  $A_n$ 
14:  end if
15: end while

```

Algorithm 2 Server Protocol

Require: List of tuples containing IP-port pairs, C
Require: Current task label, L
Require: Current task reward, R
Require: Constant identifier for query, I

- 1: Listen for incoming connections
- 2: Construct query as array $\{C_{ip}, C_{port}, I, L, R\}$
- 3: Serialize query into a byte string
- 4: **for** (C_{ip}, C_{port}) in C **do**
- 5: Sleep for 0.2 seconds
- 6: **if** Length of response list > 0 **then**
- 7: Select best agent A_n using condition $A_n = \{a \in C : 0.9R_n > R_x\}$
- 8: Send mask request to A_n for $ip_n, port_n$
- 9: **else**
- 10: Pass
- 11: **end if**
- 12: **end for**

Hardware specifications		
	Linux GPU Server	Windows GPU Machine
Architecture	x86_64	x86_64
CPU(s)	128	32
Thread(s) per core	2	2
Core(s) per socket	64	16
CPU Model	AMD EPYC 7713P	AMD Threadripper Pro 3955WX
L1d cache	2 MiB	96 Kib
L1i cache	2 MiB	N/A
L2 cache	32 MiB	512 Kib
L3 cache	256 MiB	64 Mib
CPU MHz	1700.000	4000.000
CPU Max MHz	3720.7029	4300.000
CPU Min MHz	1500.0000	3900.000
GPU 1	NVIDIA A100	NVIDIA A5000
GPU 2	NVIDIA A100	N/A
Interface	PCIe	PCIe
Memory	256 GB ECC	64 GB
GPU Memory	40 GB ECC ($\times 2$)	24 GB ECC
Base GPU Clock	765 MHz	1170
Boost GPU Clock	1410 MHz	1695
FP32 TFLOPS	156	55.55
FP16 TFLOPS	312	111.1

Table 2: Hardware specifications for the GPU servers used to run the experiments presented in this paper. Experiments featured in Fig. 2, 4 and 8 were run using the Linux-based GPU server while experiments featured in Fig. 3 was run using the Windows-based machine. Experiments were run using FP16 and FP32 configurations.

ML Packages	
Package	Version
Python	3.9.13
PyTorch	1.13.0
Gym API	0.24.0
Gym-MiniGrid	1.1.0
Gym-CTgraph	0.1
Tensorboard	2.9.1
TensorboardX	2.5.1
Pandas	1.4.3
NumPy	1.23.1
CUDA	11.7.1
cpython	0.29.30
gcc	12.2.0

Table 3: Core packages used in a single L2D2-C agent. The core ML is based on PyTorch with CUDA. OpenAI Gym is used as the interface between the model and the reinforcement learning environment. TCP/IP communication is implemented using the built-in socket module in Python.

Hyper-Parameters	
learning rate	0.00015
cl preservation	supermasks
number of workers	1
optimiser function	RMSprop
discount rate	0.99
use general advantage estimator (gae)	True
gae_tau	0.99
entropy weight	0.00015
rollout length	512
optimisation epochs	8
number of mini-batches	64
ppo ratio clip	0.1
iteration log interval	1
gradient clip	5
max steps (per task)	42752
evaluation episodes	25
require task label	True
backbone network seed	9157
experiment seeds	958, 959, 960, 961, 962

Table 4: Agent hyper-parameters. These parameters were used across all the single-server experiments shown in Fig. 3. These hyper-parameters are consistent across all experiments shown in the figure. IMPALA, DDPPO, and PPO experiments were run using the Ray API (Anyscale-Inc, 2023).

Hyper-Parameter	CT-graph	Minigrid	CT-graph Drop conn.
learning rate	0.00015	0.00015	0.00015
cl preservation	supermasks	supermasks	supermasks
number of workers	4	4	4
optimiser function	RMSprop	RMSprop	RMSprop
discount rate	0.99	0.99	0.99
use general advantage estimator (gae)	True	True	True
gae_tau	0.99	0.99	0.99
entropy weight	0.00015	0.1	0.00015
rollout length	128	128	320
optimisation epochs	8	8	8
number of mini-batches	64	64	64
ppo ratio clip	0.1	0.1	0.1
iteration log interval	1	1	1
gradient clip	5	5	5
max steps (per task)	12800	102400	12800
evaluation episodes	25	5	25
require task label	True	True	True
backbone network seed	9157	9157	9157

Table 5: Agent hyper-parameters. These parameters were used across all the single-server experiments shown in Fig. 2 and 4. Fig. 4 specifically uses the connection drops version of the CT-graph hyper-parameters.

Agents	seed 1	seed 2	seed 3	seed 4	seed 5
1	9158	6302	8946	5036	7687
2	9159	1902	4693	8814	1029
3	9160	4446	3519	2851	4641
4	9161	9575	6652	4719	7122
5	9162	1954	6613	3672	6470
6	9163	3972	4197	3523	5614
7	9164	5761	8640	6978	4687
8	9165	8004	6738	1399	8024
9	9166	3993	8248	5952	3218
10	9167	6553	7423	9744	8224
11	9168	8805	9725	3633	2047
12	9169	5066	8760	4131	6219
0	9157	9802	9822	2211	1911
13	9157	9802	9822	2211	1911

Table 6: Seeds used across each agent. These remained consistent throughout each experiment. Global evaluation agents correspond to agent 0 (evaluate the entire collective), and local evaluation agents correspond to agent 13 (evaluate a singular agent from the collective). A single L2D2-C agent corresponds to agent 1, and L2D2-C collectives are comprised of subsets of these agents (i.e., L2D2-C 4 agents consist of agents 1-4 while L2D2-C 12 agents consist of agents 1-12).

Hyper-Parameter	Depth 2 16-task CT-graph
general seed	3
tree depth	2
branching factor	2
wait probability	0.0
high reward value	1.0
fail reward value	0
stochastic sampling	false
reward standard deviation	0.1
min static reward episodes	0
max static reward episodes	0
reward distribution	needle in haystack
MDP decision states	true
MDP wait states	true
wait states	[2, 8]
decision states	[9, 11]
graph ends	[12, 15]
image dataset seed(s)	1-4
1D format	false
image dataset seed	1
number of images	16
noise on images on read	0
small rotation on read	1

Table 7: CT-graph environment parameters. We use 5 seed variations of a depth 2 4-task CT-graph to achieve the 16-task CT-graph for our experiments. We use the depth 2 16-task CT-graph for Fig. 2 (A, C, D, G, H) and 4.

Hyper-Parameter	3-task Minigrid
tasks	MiniGrid-SimpleCrossingS9N1-v0, MiniGrid-SimpleCrossingS9N2-v0, MiniGrid-SimpleCrossingS9N3-v0
one hot	true
label dimensions	3
action dimensions	3
seeds	860, 860, 860

Table 8: Minigrid environment parameters. Each seed corresponds to task IDs 0-2 respectively. We use the 3-task Minigrid for Fig. 2 (B, E, F, I, J).

agents/tasks	Time to X% performance			Performance at X% time		
	20%	50%	70%	20%	50%	70%
2/2 (L2D2-C)	1.7	27.9	33.1	6.6	100	100
1/2 (singleLLAgent)	4.1	33	57	0	50	48.3
4/4 (L2D2-C)	10.2	12.5	16.3	71.2	97.5	100.0
1/4 (singleLLAgent)	11.1	45.8	76.0	22.5	43.1	50.0
8/8 (L2D2-C)	5.4	6.9	16.3	72.7	98.7	98.5
1/8 (singleLLAgent)	22.4	55.8	79.9	12.9	32.7	57.0
16/16 (L2D2-C)	2.9	4.5	8.6	100.0	100.0	100.0
1/16 (singleLLAgent)	24.6	57.5	76.9	12.5	40.6	56.2
32/32 (L2D2-C)	1.8	2.7	4.5	96.8	100.0	96.8
1/32 (singleLLAgent)	22.8	60.5	79.2	12.5	37.5	54.6
4/16 (L2D2-C)	5.1	18.8	28.0	55.4	98.7	100
1/16 (singleLLAgent)	24.5	55.2	76.2	12.5	37.5	59.6

Table 9: Time to X% performance (in %) and performance (in %) at X% time (with total time measured as the time required to reach 95% performance). Due to noise affecting evaluations with few tasks, smoothing was applied: percentages are more representative for experiments with 8 or more tasks. L2D2-C 16/16 and 32/32 (in bold) were the fastest vs the single agent.

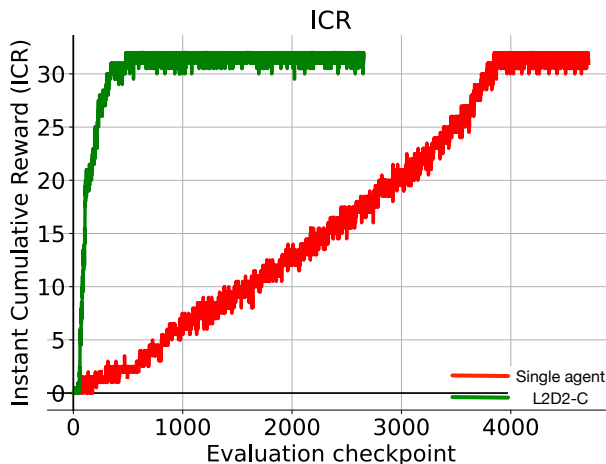


Figure 8: Instant cumulative reward (ICR) for the 32-agents/32-tasks experiment.

C ADDITIONAL RESULTS

Table 9 reports additional metrics for single-seed experiments in which the agents/tasks ratio was used: 2/2, 4/4, 8/8, 16/16, 32/32, and 4/16. The instant cumulative reward (ICR) for the 32/32 experiment is reported in Fig. 8.

Fig. 9 shows the performance of an L2D2-C system across two locations in different countries. This is a proof-of-concept that the system can work across locations and more statistical analysis and experiments are required. However, it indicates that four agents across two locations do not appear slower than 4 agents in the same location.

Scatter plots showing the different types of messages exchanged by all agents during 4 runs with different probabilities of connection drop are shown in Fig. 10.

D ENVIRONMENTS

D.1 CT-GRAPH

In the configurable tree graph (CT-graph) (Soltoggio et al., 2019; 2023) the environment is implemented as a graph, where each node is a state represented as a 12×12 grayscale image. As shown in Fig. 11, the node types are start (H), wait (W), decision (D), end/leaf (E), and fail (F). The goal of an agent is to navigate from the home state to one of the end states designated as the goal where the agent receives a reward of 1. To navigate to any end of the graph, the agent

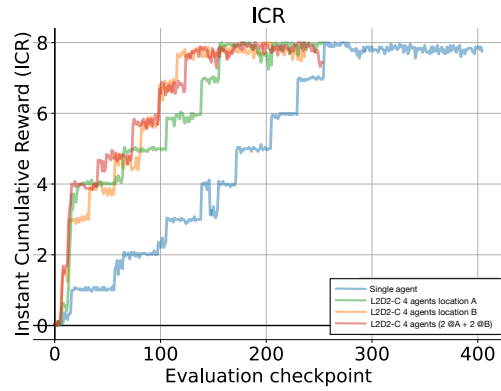


Figure 9: An L2D2-C system run across two locations in two different countries. The system across two locations does not appear to be slower than the system that ran in the same location.

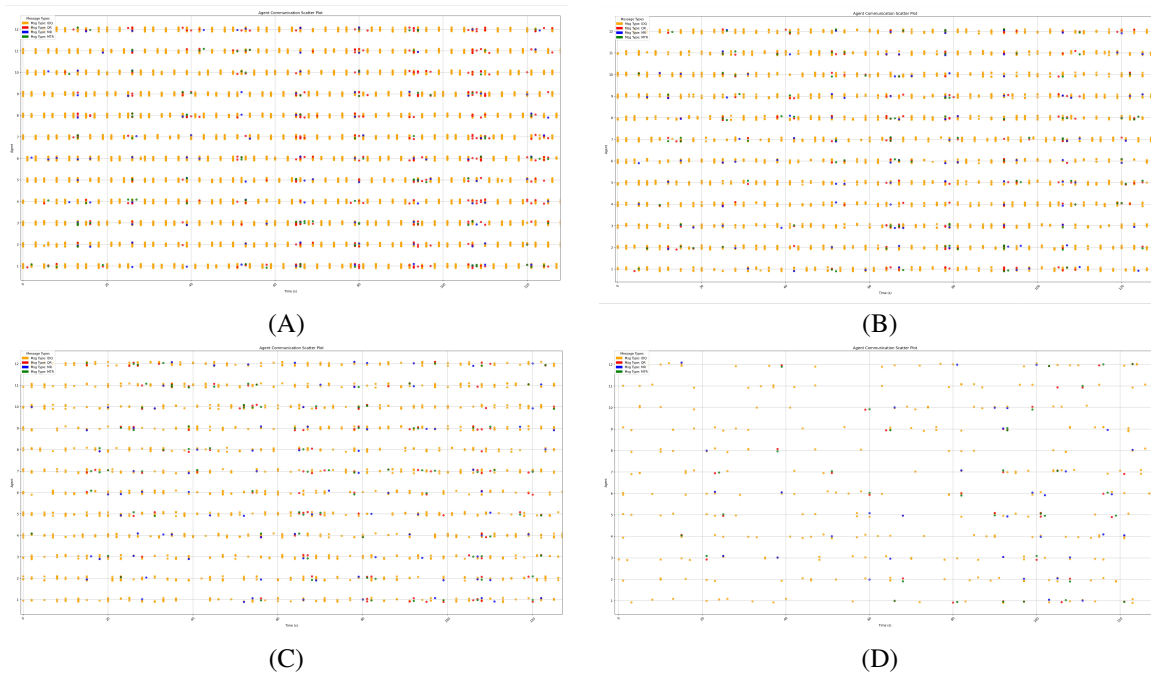


Figure 10: Different message types exchanged by 12 agents during runs with different connection drop probabilities. For better visibility, a time window of 120 seconds only is shown. (A) Full connection. (B) 50% connection drop probability. (C) 75% connection drop probability. (D) 95% connection drop probability.

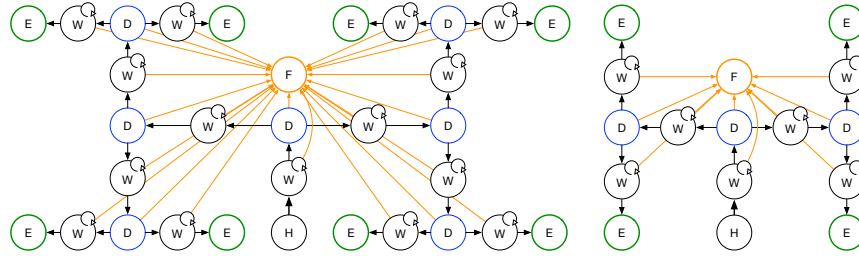


Figure 11: Examples of CT-graph environments. States are: home (H), wait (W), decision (D), end (E), and fail (F). Three actions at W and D nodes determine the next state. (Left) A depth-3 graph with three sequential decision states (D) and reward probability $1/3^7 = 1/2187$ reward/episodes. (Right) A depth-2 graph with 4 leaf states.

is required to perform action a_0 when in W states, and any other action a_i with $i > 0$ in D states. The challenge is any policy that does not follow the above criteria will lead the agent to the fail state and terminate the episode without reward.

Two parameters in the CT-graph are the branching factor b that determines how many sub-trees stem from each D state, and the depth d that determines how many subsequent branching node D occur before a leaf node. By setting these two parameters, different instances of the CT-graph can be created with a measurable size of the search space and reward sparsity.

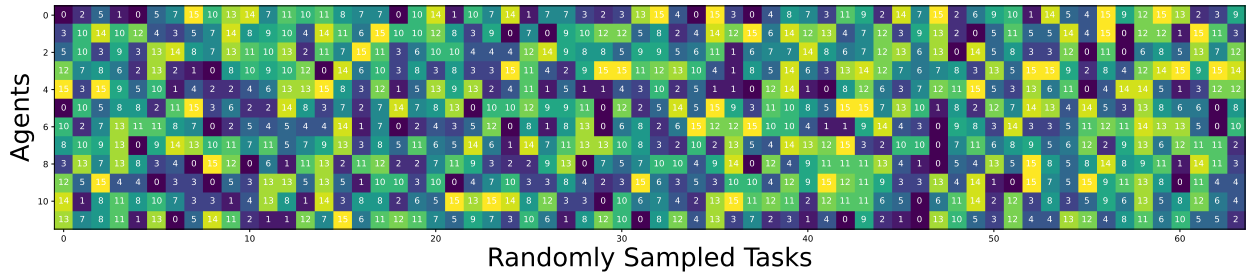
In a typical execution with the CT-graph, tasks are sampled randomly. Fig. 12 shows a typical curriculum that was visualized and analyzed.

D.2 MINIGRID

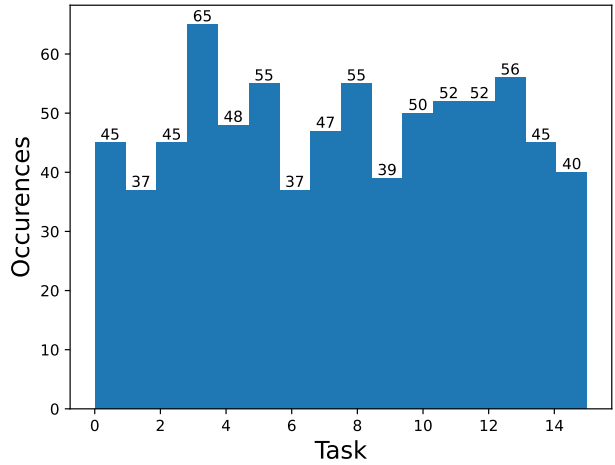
The Minigrid environment (Chevalier-Boisvert et al., 2018) is a grid world where an agent is required to navigate to the goal location. There are pre-defined grid worlds with sub-variants according to the random seed. A tensor of shape $7 \times 7 \times 3$ is used as input. The reward is defined as

$$goal_reward = 1 - 0.9 \times \frac{es}{ms} \quad , \quad (3)$$

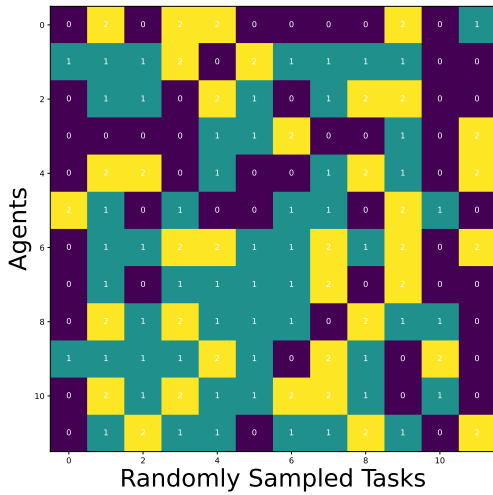
where es is the number of steps taken to navigate to the goal and ms is the maximum number of steps in an episode. The environment variations employed are: SimpleCrossingS9N1, SimpleCrossingS9N2, SimpleCrossingS9N3. Fig. 13 shows the five grid worlds used.




(A)



(B)



(C)



(D)

Figure 12: Distribution of randomly sampled tasks. (A) 12 sequences of tasks (one for each agent, plotted as 12 rows) are visualized with each task having a different color. The horizontal dimension is the time, showing here that there are 64 task changes. This task curriculum consists of tasks randomly sampled from a 16-task CT-graph. (B) The number of times each task is encountered across all learning time and agents. (C) 12 sequences of tasks are visualized for 12 task changes for 3 Minigrad task variations. (D) The number of times each task is encountered across all learning time and agents. For experiments shown in Fig. 4 we use a modified curriculum in which the first agent uses a sequential repeating curriculum i.e., 0 15 for 4 cycles. The first agent also acts as the target of the local evaluation agent featured in figure 4 and so it is a reflection of the learning dynamics of the first agent when it is connected to the rest of the collective with varying amounts of connection drops.




Figure 13: Visual representation of the 3 tasks in the *MG3* curriculum. From left to right, one variant of each class: SimpleCrossingS9N1, SimpleCrossingS9N2, SimpleCrossingS9N3. The agent (red) must navigate the environment with impassable walls (gray) and reach the goal state (green). The agent can only observe a limited portion of the environment, indicated by the highlighted squares in the images.

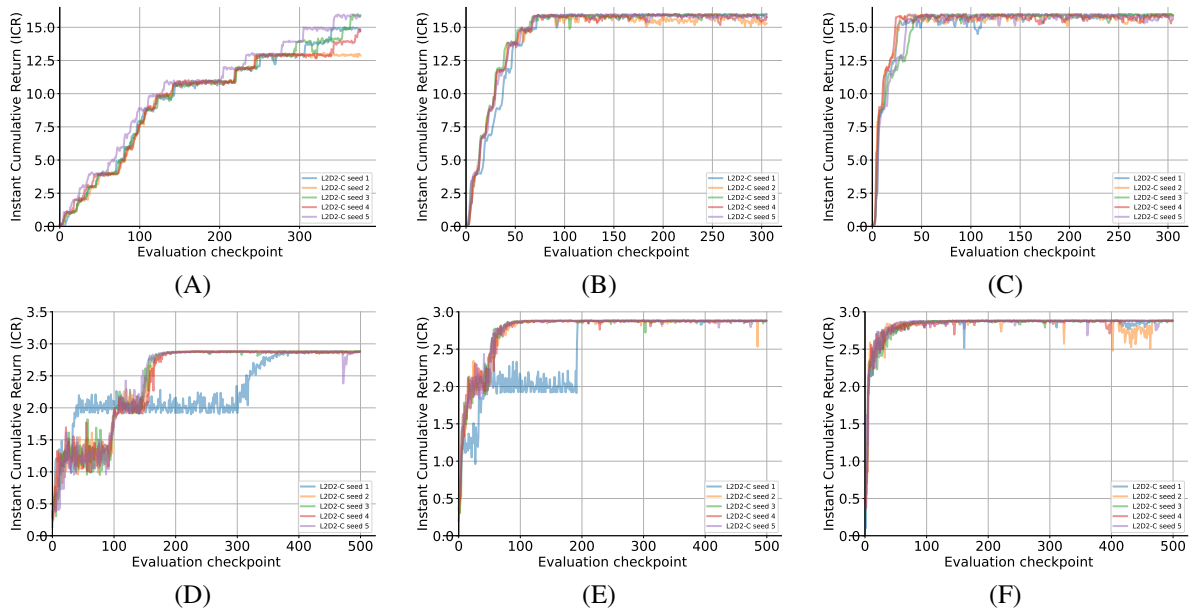


Figure 14: Individual seed runs of L2D2-C on CT-graph and Minigrad from Fig. 2. A total of 5 seeds were used for each experiment. The graphs show configurations of 1, 4, and 12 agents tested on the CT-graph environment (A, B, C) and similarly on the Minigrad environment (D, E, F).