nature machine intelligence

Perspective

https://doi.org/10.1038/s42256-024-00800-2

A collective AI via lifelong learning and sharing at the edge

Received: 30 April 2023

Accepted: 24 January 2024

Published online: 22 March 2024

Check for updates

Andrea Soltoggio[1] ✉, Eseoghene Ben-Iwhiwhu[1], Vladimir Braverman[2], Eric Eaton[3], Benjamin Epstein[4], Yunhao Ge[5], Lucy Halperin[6], Jonathan How[6], Laurent Itti[5], Michael A. Jacobs[2,7,8], Pavan Kantharaju[9], Long Le[3], Steven Lee[10], Xinran Liu[11], Sildomar T. Monteiro[10,12], David Musliner[9], Saptarshi Nath[1], Priyadarshini Panda[13], Christos Peridis[1], Hamed Pirsiavash[14], Vishwa Parekh[15], Kaushik Roy[16], Shahaf Shperberg[17], Hava T. Siegelmann[18], Peter Stone[19,20], Kyle Vedder[3], Jingfeng Wu[21], Lin Yang[22], Guangyao Zheng[2] & Soheil Kolouri[11]

One vision of a future artificial intelligence (AI) is where many separate units can learn independently over a lifetime and share their knowledge with each other. The synergy between lifelong learning and sharing has the potential to create a society of AI systems, as each individual unit can contribute to and benefit from the collective knowledge. Essential to this vision are the abilities to learn multiple skills incrementally during a lifetime, to exchange knowledge among units via a common language, to use both local data and communication to learn, and to rely on edge devices to host the necessary decentralized computation and data. The result is a network of agents that can quickly respond to and learn new tasks, that collectively hold more knowledge than a single agent and that can extend current knowledge in more diverse ways than a single agent. Open research questions include when and what knowledge should be shared to maximize both the rate of learning and the long-term learning performance. Here we review recent machine learning advances converging towards creating a collective machine-learned intelligence. We propose that the convergence of such scientific and technological advances will lead to the emergence of new types of scalable, resilient and sustainable AI systems.

Progress in science, technology and other fields of knowledge is largely possible due to the ability of individual humans to build on discoveries and knowledge from other individuals[1]. New challenges are faced by leveraging knowledge accumulated over time and transferred from individual to individual. Although no single person can possess all knowledge and intelligence, collectively and as a species, we have the remarkable ability to acquire knowledge from others, adapt it, extend it further and explore different ideas and methods thanks to different objectives and predispositions while also maintaining agency and individuality[2]. A natural question is whether progress in AI could see a similar structure emerging from a society of AI agents. A new type of AI may be based on a collective effort of many individual entities that, similarly to humans, can acquire knowledge from other entities, add to it, and maintain a diverse and decentralized structure.

Multi-agent and distributed AI systems have been studied for decades[3–5], albeit often with a focus on cooperation to solve a single task[6] or competition in game theory studies[7]. In this paper, we focus specifically on the attempts to share machine-learned knowledge among a collective of agents with the aim of enhancing the performance of each individual agent. To this end, agents share their

A full list of affiliations appears at the end of the paper. ✉e-mail: a.soltoggio@lboro.ac.uk

knowledge and do so not to deceive or undermine the performance of other agents. Each individual tries to improve their own individual performance by exploiting both their own experience and knowledge acquired via other agents. Under these assumptions, a collective of agents may result in beneficial interactions among agents that can maximize their performance thanks to knowledge reuse and sharing.

The perspective of such a collective AI is becoming more realistic thanks to recent advances in fields such as lifelong learning (LL)[8–12], lifelong reinforcement learning[13,14], federated learning[15], distributed and multi-agent systems[16], and edge computing[17,18]. Although these fields address specific challenges, integrating advances from all such areas may result in a significant step towards collective machine learning (ML) ecosystems. The efforts to unify such fields have questions for research. What information should be exchanged among agents and when is it best to share? If thousands or millions of agents learn an equivalent number of tasks, how does a single agent determine what information is relevant to them? How is new information integrated into one agent's knowledge to expand its capabilities without forgetting? What hardware platforms will enable a fully decentralized system capable of performing and learning individually while also communicating[19]?

Efforts to answer the above questions are contributing to producing a new type of decentralized AI with properties resembling some aspects of learning in humans. One benefit of a distributed and decentralized society of learning agents is the robustness and resilience to failure or adversity, a concept largely explored in various network applications[20], including the Internet, peer-to-peer networks and blockchain. As opposed to large and centralized models, if knowledge is acquired and maintained on smaller edge units capable of sharing, the loss of some agents may result in a limited loss of the collective knowledge, which can then be acquired by newly spawned agents. Failures or mistakes of such agents will also result in less catastrophic consequences than the failure of one single central model that is responsible for controlling large systems. The acquisition of knowledge via other agents also implies that experiences, particularly negative or dangerous ones[21], need not be repeated. Effectively, sharing endows more power to the search as parallel executions will reveal properties of large search spaces. Many agents, by sharing knowledge, can sum their efforts to achieve collectively faster and more complete learning on a large amount of data, as is being showcased by federated learning approaches[22,23] and distributed ML[24]. With a worldwide increase of computational demands for training ML models, a collective of AI agents that can reuse and share knowledge may be a solution to better scalability and reduced energy demands[25,26]. Finally, learning a large variety of tasks and integrating them into each agent's individual models may lead to a more structured and explainable organization of knowledge. The end result of the above properties is an increased ability of artificial agents to expand their knowledge toward open-ended applications, efficiently sharing and reusing knowledge.

This Perspective is organized as follows. First, we highlight the research fields that contribute to the emergence of lifelong learning and sharing systems before describing their main objectives, capabilities and technical aspects. Then, we review recently developed approaches that showcase the feasibility and potential of the methods. Finally, we describe relevant application areas before discussing outstanding challenges and opportunities.

Constituent fields of a lifelong learning collective

The idea of a lifelong learning collective emerges from the integration and synergies of the fields described in Fig. 1. We refer to such a convergence of concepts as shared experience lifelong learning (ShELL)[19]. The contributing fields in Fig. 1 are not isolated efforts and often extend along the connecting edges in the research space with approaches that combine principles from different fields. Although the absolute importance and the precise boundaries of each of the listed fields may be subjective, each of them contributes to advances that are relevant to the vision presented in this Perspective.

Key aspects of integrating knowledge over a lifetime

Lifelong ML (LL) refers to a set of methods that allow an algorithm to learn from a continuous stream of data and different tasks over a prolonged period of time[8–10,27]. Such methods have also been referred to as incremental or continual learning[28]. The first solutions emerged with a focus on preventing catastrophic forgetting[29,30]. In a typical ML algorithm, a model improves its performance while it trains on a curated set of identical and independently distributed (IID) data, implying that all learning examples must be present from the start and describe the final data distribution well. Continual or incremental learning has focused on methods that can train a model when different data distributions are seen sequentially[31–33] by adding various methods to mitigate catastrophic forgetting. These approaches have clear advantages when datasets are incomplete or change with different distributions over time.

Crucially, the ability to learn from sequential data without forgetting created the possibility to remove the established distinction between a training phase, when the model is plastic, and a deployment phase when the model is frozen and therefore unable to improve further. With the objective to create always-learning algorithms that do not have separate training and deployment phases[34], can improve over time and can deal with surprise events and unexpected data, the Defense Advanced Research Projects Agency (DARPA) funded the Lifelong Learning Machine programme[35]. This programme investigated both biological inspiring principles[11] and core LL algorithms, improving the understanding of algorithmic properties and metrics for evaluation[36,37]. Currently, as a result of the evolution of the field, LL is not limited to reducing catastrophic forgetting but includes learning dynamics such as the ability to exploit previous knowledge to learn new tasks[38–40], use experience in one task to improve performance on other related tasks, generalize knowledge over many tasks[41] and by the synergy of those abilities, produce a substantially more capable and open-ended learning system.

LL algorithms promise to be more flexible than traditional ML but generally involve additional computation, memory or storage, and hyper-parameters that require careful consideration. The first challenge is that additional computational and memory requirements vary according to different families of LL approaches (for example, replay, regularization or parameter isolation methods[32]): the exact overhead in terms of memory and computation, and their suitability to edge deployment has only recently become a subject of investigation[42–44]. A second challenge is that LL algorithms are designed to learn continuously while they are deployed: as a consequence, the computational cost of learning may no longer be offloaded to remote servers. Inference and learning need to happen at the same time during deployment, often resulting in the execution of parallel operations that need to be skillfully coordinated and executed locally. Finally, the evaluation of LL systems is not as straightforward as in traditional ML where testing is generally performed after training and before deployment, and evaluation metrics are well established. In LL, evaluation is required during deployment on multiple tasks and includes additional metrics that assess levels of forgetting, the ability to reuse knowledge (forward and backward transfer) and others[36,37,45–47].

Despite the challenges described above, real-world applications of ML demand increasingly more adaptive systems that can continuously improve during deployment, face a large variety of new and diverse tasks[48], and mimic lifelong learning properties of biological systems[11,49]. In the vision that we present in this paper, LL algorithms are an essential building block for an AI collective: exchanging knowledge among agents is made effective by the capability of integrating knowledge learned at different times, on different tasks and by different agents.

https://doi.org/10.1038/s42256-024-00800-2

a

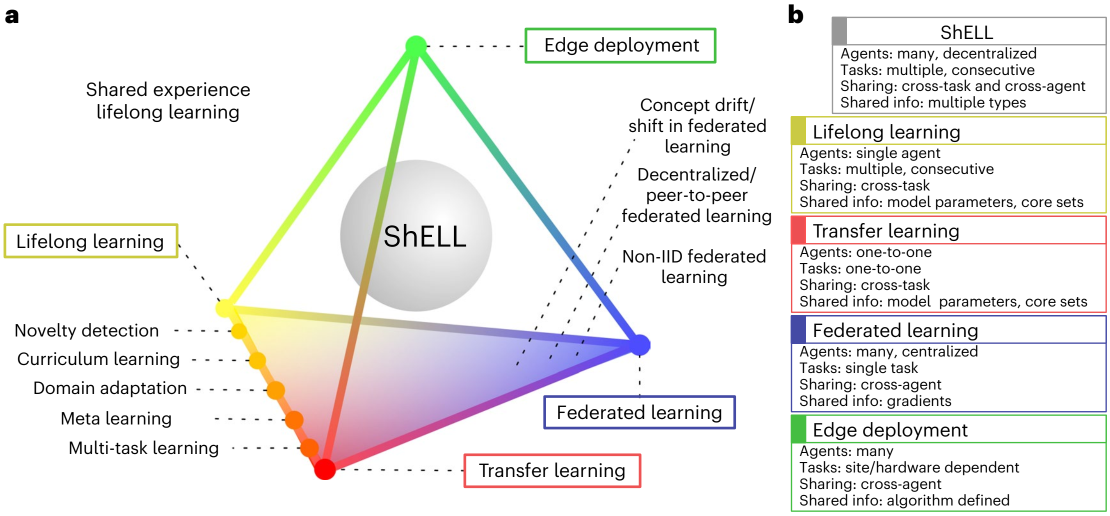

c

Algorithmic approach	Agents	Consecutive tasks	Multiple tasks	Knowledge reuse	Non-IID data for one task	Distributed	Decentralized	Sharing	On-device
Transfer learning	1	✓	✓	✓	○	○	○	✓	○
Multi-task learning	1	○	✓	▲	○	○	○	○	○
Lifelong learning	1	✓	✓	✓	✓	▲	○	○	▲
Federated learning	Many	○	▲	○	▲	✓	▲	✓	✓
ShELL	Many	✓	✓	✓	✓	✓	✓	✓	✓

✓ Core
▲ Emerging
○ Unusual

Fig. 1 | Research fields contributing to ShELL. a, ShELL emerges at the intersection of the listed fields by integrating concepts that are required to produce lifelong learning and sharing. Shared experience is interpreted as a general information-sharing mechanism that can involve a wide range of categories such as data, context, parameters, gradient updates, representations and algorithms. The listed fields have originated with a different focus on the number of agents, number of tasks, sharing modalities and objectives. As such research fields evolve, they expand along the edges of this abstract space to include more capabilities. Algorithmic approaches listed on the lower plane require integration along the vertical dimension when deployed on edge computing devices. **b**, A summary of the main properties defining agents, tasks, sharing and shared information for the main fields listed. **c**, A non-exhaustive table of a few related ML approaches and properties shows how ShELL borrows and combines methods from related areas.

Distributed machine learning with federated learning

The ability to perform ML on a large number of distributed edge devices has been extensively investigated in the field of federated learning[15,23,50,51]. The main aims when performing training on edge devices[52,53] are to reduce data transfer and maintain privacy. To do so, parameter updates are computed locally using local data by each client. Only parameter updates (for example, gradients) are shared with a central server that aggregates updates from a large number of clients. Federated learning algorithms take into account power management, limited computational resources and memory. Commercial applications have taken advantage of this technology to comply with privacy and regulatory issues.

Recently, federated learning studies have started to consider LL-related issues, such as non-IID data across locations and time[54–56], thus identifying the need to prevent catastrophic forgetting[57,58]. In particular, in domains where personalization and geographical variations are important, the addition of LL, multi-task or meta-learning algorithms to federated learning has become a priority[59,60]. Although not a core focus of federated learning, by engaging with datasets and scenarios that require multi-task learning and LL (Fig. 1), federated learning contributes algorithmic novelty and experimentation to the fields of LL and sharing. Most federated learning approaches, although distributed, update a single central model, however decentralized approaches have been reported[61].

Exploiting task similarity and reusing knowledge

The idea of reusing knowledge in ML has been extensively investigated in many related research areas. Reusing previously acquired knowledge when learning a new task is key to transfer learning[62–64]. The assumption is that a degree of similarity between the old and the new task will make previously acquired knowledge useful. One popular application of transfer learning is the use of models trained on large datasets as a starting point to train on personalized problems with small datasets. This approach has obvious benefits in real-world applications of ML, including recently developed methods for large language models (LLMs) namely parameter efficient fine-tuning[65,66]. When the differences in distributions between tasks can be estimated, domain adaptation provides methods to adapt the models accordingly[67–70].

The same assumption that task similarities can be exploited motivates multi-task learning[71–74]. Instead of learning different models for different tasks, the idea is that it is more efficient to learn one model that can solve a set of tasks because common features among tasks need not be learned multiple times. A similar idea is also exploited in meta-learning[75], where the search objective is to find a model that can be quickly adapted to a number of tasks. Meta-learning was also shown to be successful in a decentralized framework[76] and in combination with continual learning[77]. The concept that specific previous knowledge is beneficial to learning new tasks is also exploited in curriculum learning[78,79], where the particular sequence and complexity of the tasks are designed to facilitate learning. Finally, reusing knowledge from a domain to solve a new problem is also central in zero-shot learning[80–84]. None of the above approaches would be beneficial if all tasks were uncorrelated. Understanding similarities among datasets[85,86] and tasks is, therefore, an emerging field that could empower knowledge reuse, help knowledge organization, detect exceptions or uncertainty[87–89], and learn without an oracle[86,90].

Given the considerable focus on knowledge reuse in the research areas described above, it comes as no surprise that the same principles are essential in ShELL. The reuse can occur within a single agent that uses past knowledge in an LL algorithm to learn new tasks, and across agents, where knowledge acquired by other agents can be shared and benefit the agents collectively. The assessment of task similarities[86] is also crucial when deciding what specific knowledge has to be transferred from one agent to another.

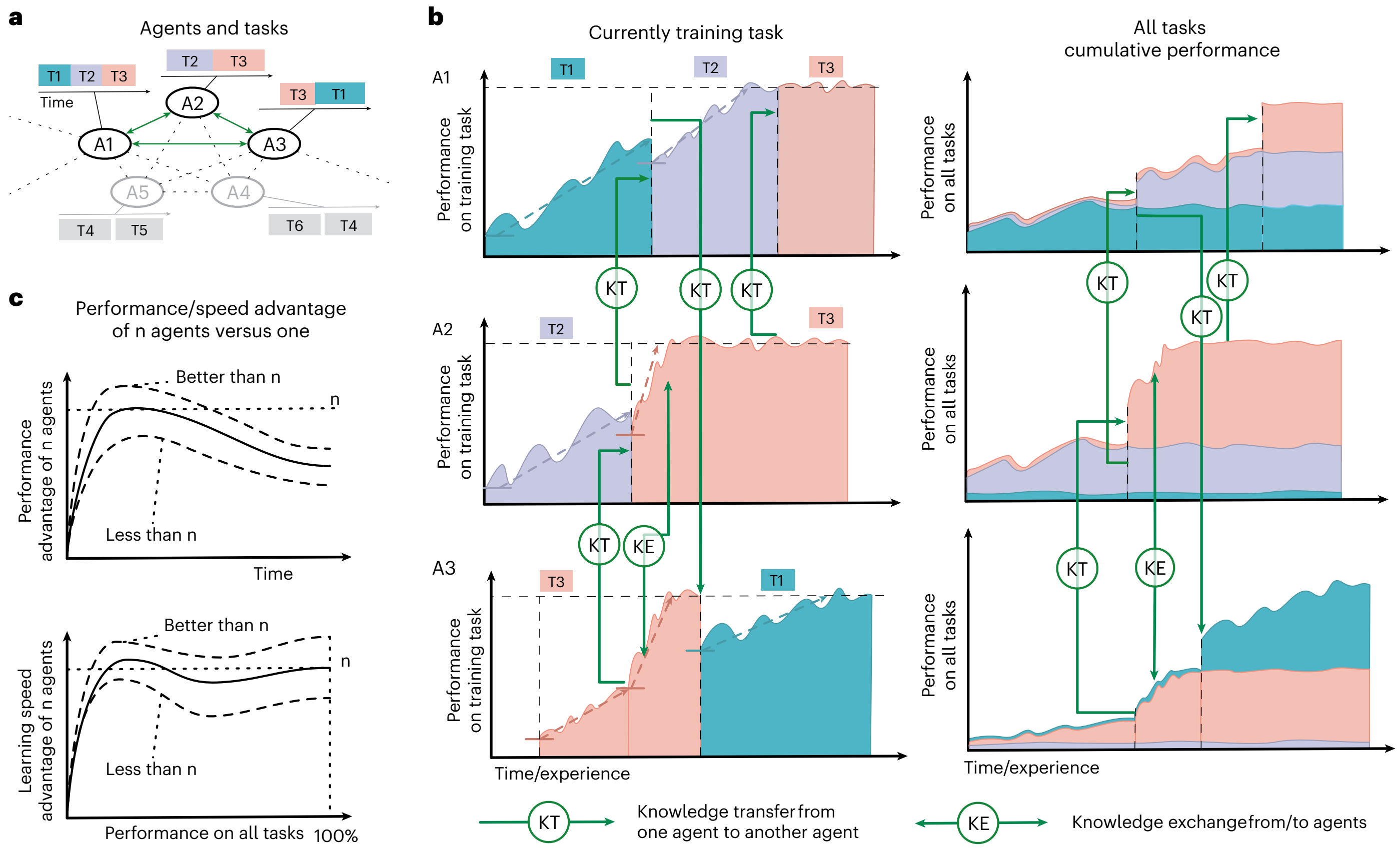

Fig. 2 | Desired learning behaviour for an on-knowledge-demand ShELL system. a, Illustration of three connected fictional agents learning three tasks in a network with multiple agents and tasks. **b**, Illustration of the desired performance of each agent on its learning tasks (left) and all tasks (right). Agent 1 (A1, top) starts on task 1 (T1) and later continues learning on T2 from the performance level that was previously reached by A2 thanks to a knowledge transfer (KT) operation. Furthermore, when A1 engages with T3, it obtains an already optimized policy from A2, and hence immediately maximizes the performance on that task. A2 and A3 engage in knowledge exchange (KE) when learning T3 at the same time, leading to faster, synergistic learning. The right plots show the performance on all tasks: agents are able to retain knowledge (LL) while both learning from their data and acquiring knowledge from other agents. **c**, Performance gain of *n* agents versus one agent: the advantage can be computed as a ratio of performance indices for the collective versus the single agent. Compared with the single agent, *n* agents can be expected to perform better (top) and be faster (bottom). The performance advantage (top graph) might decrease over time when all tasks are learned by the collective (if the number of tasks and their distributions do not change), as the single agent continues to learn tasks. These fictional graphs attempt to summarize trends from experimental evidence observed in ShELL studies reported in the 'Emerging approaches to lifelong learning and sharing' section and in the Supplementary Information.

Synergy between lifelong learning and sharing

The processes that allow for the integration of knowledge both from an agent's own experience, via LL, and from communication, via the acquisition of external knowledge, can be combined to produce a collective of ML agents with the following objectives:

(1) Collectively learn multiple tasks faster than individual agents thanks to sharing information and pooling resources.
(2) Continually adapt and increase knowledge over a lifetime and across a population of agents.
(3) Exploit decentralized agents to learn independently and to hold diverse information or policies leading to better robustness and diversity in the solutions.

To achieve these objectives, each agent requires the algorithmic and technical abilities to:

(1) Integrate knowledge (over a lifetime) learned from both its own sensors and data and from communication with other agents.
(2) Reuse knowledge from self-learning or from other agents to learn new tasks at a faster rate.
(3) Request or send specific machine-learned knowledge from or to other agents.
(4) Self-organize learning and communication among agents without a central authority.

A system that implements these four abilities is expected to demonstrate performance dynamics to reflect the three objectives. In addition to established ML metrics[91] and recently developed LL metrics[36,37,45–47,92], ShELL-specific metrics may include indicators that convey the advantage of sharing knowledge. Figure 2a illustrates three fictional agents interacting while learning different tasks consecutively. The desired learning dynamics of the agents are illustrated in Fig. 2b. The learning speed on a single task (Fig. 2b, left) is boosted by initiating learning from a non-zero performance level, thanks to knowledge transfer, or by more than one agent collaborating on learning one task via knowledge exchange. When summing the performance on all tasks (Fig. 2b, right), the ability to integrate knowledge without forgetting results in near-monotonic growth of the overall performance.

Considering a single LL agent as a baseline, two metrics can be defined to measure objective 1: (1) how much performance gains can *n* agents achieve within a determined time with respect to a single agent; (2) how much faster can *n* agents achieve a predetermined performance level with respect to the single agent[19]. The expectations are that *n* agents can learn faster or better because more resources are deployed in parallel, and communication enables each agent to access

https://doi.org/10.1038/s42256-024-00800-2

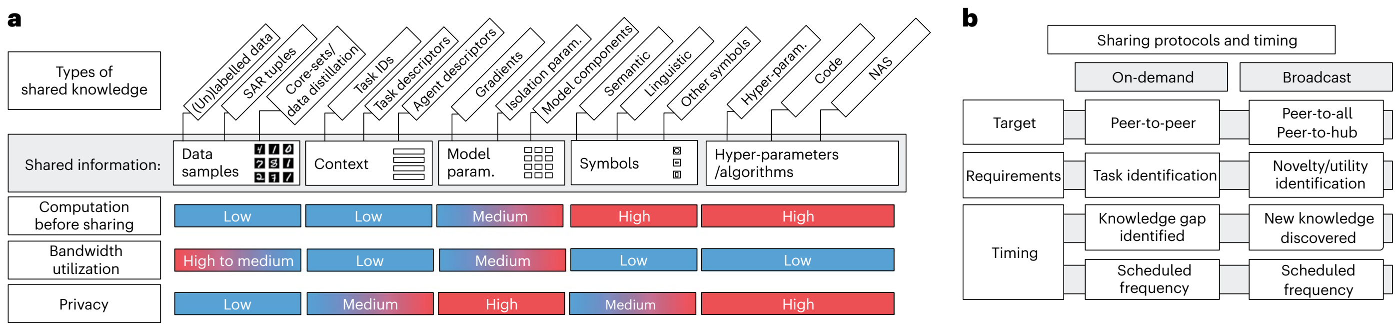

Fig. 3 | Types of information sharing and timing. a, Trade-offs in computation and bandwidth. Different types of data lead to different utilization of computational resources and bandwidth and to different degrees of privacy. From left to right, the computation before sharing increases, resulting in more compact knowledge representations being shared. The implication is that the more is learned at the edge (near the data) the less bandwidth is required to communicate among devices. **b**, Sharing protocols and timing. On-demand and broadcast result in different communication targets, requirements and timing. All protocols can have a scheduled agent-dependent frequency of communication. All communication is asynchronous and not centrally coordinated.

the collective's knowledge, albeit possibly with loss in efficiency or time delay. Figure 2c illustrates possible trends in the advantage of the collective with respect to a single agent in relation to better performance (top) and an increase in learning speed (bottom). Although such metrics have been proposed previously[19] for measuring objective 1, ShELL metrics are currently under development, particularly with additional metrics required to capture how well objectives 2 and 3 are achieved. Further considerations on the expected performance of ShELL are reported in the Supplementary Information.

Types and timing of knowledge sharing

Agents may share different types of information with different impacts on how much computation is required before sharing, how much bandwidth is used and whether privacy is preserved (Fig. 3a).

Agents may share curated data that best describes a task. Data exchange can be model agnostic but is generally paradigm specific: that is, labelled data for supervised learning, unlabelled data for unsupervised learning and SAR tuples (state, action, reward) for reinforcement learning. Data can be necessary when performing knowledge distillation across models[93]. One advantage of sharing data is that agents could make representative data available, that is, core datasets, that best capture the distribution of the task, or capture exceptions or anomalies. Disadvantages of exchanging data are bandwidth and memory requirements for high-dimensional data points such as images, and lack of privacy across agents or locations. This could limit applications in bandwidth and memory-constrained scenarios.

Agents learning different tasks are required to obtain or infer a task label or context information, which may be shared[94]. In addition, agents may share information such as what task they are currently learning, what tasks they are capable of solving, and with what performance. Novelty detection[87,89,95] and context information[96,97] are particularly relevant for applications with many, possibly interfering, tasks.

Exchanging model parameters allows agents to share information that has already been extracted from data. A compact set of parameters can incorporate knowledge of a large amount of data. Bandwidth and memory utilization may be considerably reduced with respect to exchanging data. However, parameter transfer is generally model-specific and, therefore, less viable with heterogeneous collectives. Rather than sharing the entire model parameters, agents may also exchange partial models. This is particularly relevant when the model can be decomposed into modular components[98], each of which is both reusable and exchangeable among agents. Examples of model components in LLMs are low-rank adaptation[99], novel model reparametrization techniques[66,100] and soft prompts[101] that allow for fine-tuning of large models by only adjusting a significantly smaller set of parameters in specific parts of the model[65]. These approaches were developed to adapt large pre-trained models to specific tasks by searching a small subset of parameters, therefore decreasing computational costs and catastrophic forgetting. While they do not explicitly include LL dynamics, they are transfer learning methods that can be used to transfer task-specific knowledge in a compact form across agents that share the same LLM. In some cases, sharing components could enable heterogeneous agents, since model architectures can differ among agents, albeit with some form of model compatibility across agents[102,103]. The exchange of network sub-regions via masks[94] or other suitable subsets[102,103] is also a form of transfer of partial model parameters. This approach can result in efficient use of bandwidth in scenarios where communication is highly constrained.

The emergence of neuro-symbolic AI[104] has shown the promise of augmenting subsymbolic systems with high-level concepts that can be used for reasoning and explicit representations. The appeal of sharing symbols in ShELL systems derives from the optimization of communication and the maximization of information transfer across agents. A commonly defined language, however, is a pre-requisite among agents.

In a meta-learning process, the outcome of learning is a model that is particularly fast at learning a set of tasks with similar distributions[75]. A similar meta-learning process occurs with neural architecture search (NAS)[105,106]. Extending such ideas to ShELL, agents may share an entire learning algorithm that was the outcome of a search process, ideally compressing the most amount of knowledge into compact representations.

Timing of sharing. The timing of sharing in ShELL is asynchronous and not centrally coordinated. This is to allow devices with possibly different computation speeds or operating schedules to contribute at different rates and enter or leave the collective at any time. Scheduled or event-based messages can be sent to broadcast learned knowledge to other agents or request knowledge when a task requires it (on-demand; Fig. 3b). Broadcasting systems require criteria to decide what information should be shared and when, but have the advantage that all agents receive new knowledge as soon as it is acquired. A scalability issue may emerge as an increased number of agents could broadcast large amounts of (potentially redundant) information to all other agents. On-demand systems make better use of bandwidth and local resources, only requesting what each agent deems necessary, but require specific task knowledge definitions and could potentially miss out on useful knowledge.

LL and sharing on edge devices

The algorithmic principles outlined above are hardware agnostic. However, one aspect that makes ShELL possible is for each agent to rely on its own computation to learn in the proximity of sensors and actuators[42,43,107,108]. As a result, the efficiency of learning algorithms, power requirements and compute/memory constraints play a more

 https://doi.org/10.1038/s42256-024-00800-2

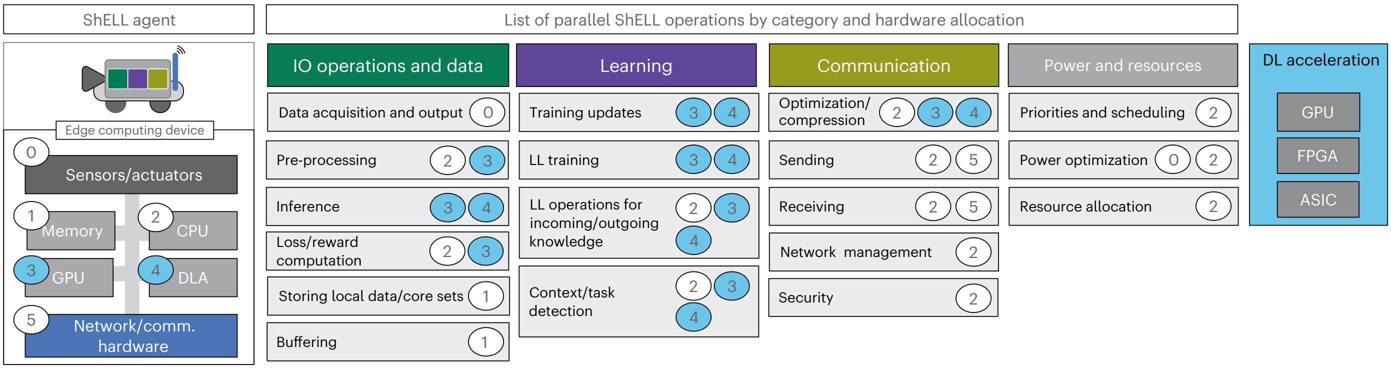

Fig. 4 | List of ShELL operations and their likely hardware allocation on a typical computer architecture on commercially available edge devices. The illustration aims to summarize the large variety of operations that ShELL requires with respect to conventional ML approaches. The four main categories of operations are shown. Specific operations for each category (grey boxes), and their hardware allocation (0 to 5) are also shown. The blue shades highlight neural-network-specific computation that mostly characterizes ShELL on edge devices. Such a computation for LL and sharing benefits from special hardware such as GPU, field-programmable gate arrays (FPGAs) and application-specific integrated circuits (ASICs) to facilitate deployment of LL algorithm to SWaP-constrained devices.

important role than in ML on large servers[109–112]. In the past decade, embedded mobile devices that include CPUs, GPUs, accelerators, network interfaces and sensors have gathered significant commercial interest. Such devices implement general-purpose computer architectures on compact mobile platforms, often including neural accelerators, and are designed with particular consideration for SWaP (size, weight and power) constraints.

ShELL objectives are not only achieved by training a model at the edge, but more broadly by combining multiple operations that can be grouped under four categories (Fig. 4): (1) input–output operations and data, (2) on-device (lifelong) learning, (3) communication and (4) power and resources management. Category 1 is typical of edge-AI devices that incorporate trained models for inference, for example, smartphones, smart cameras and autonomous vehicles[17,113,114]. The successful deployment of neural models on edge devices has led manufacturers to incorporate deep learning accelerators for inference on most portable devices[59,115] and to the optimization of deep learning algorithms[116]. Category 2 is required by the on-device model training[117,118] and LL algorithms. The requirement for deployment on small devices has lead to studies that measure and optimize computational loads of LL algorithms[18,44,47,119–121], and more generally for model training with an emphasis on efficiency, for example, via quantization[122–125]. The computational load of LL approaches is affected by the specific type of algorithm, for example, replay[126], regularization or parameter isolation methods[32,44], which could involve different use of CPU, GPU or memory[127]. For agents that operate with real-time input–output requirements, operations in category 2 need not slow down or affect the operations in category 1. The demands for computation with artificial neural networks both at the inference and learning stages (categories 1 and 2) motivate research in more powerful accelerators: in addition to GPUs and field-programmable gate arrays, application-specific integrated circuits[128] have been developed for ML tasks, including tensor processing units (TPU)[129], and emerging computational approaches such as in-memory computing[130], neuromorphic computing[131–135], spiking neural networks[135,136], integrated photonic tensor core[137–139] and others[128]. Communication (category 3) is a common category of edge devices. In ShELL, a communication module is required to package and send specific lifelong learning knowledge, extracted or managed by category 2 operations. Finally, power and resource management[140,141] (category 4) is critical in ShELL to ensure that the operations in categories 1–3 are optimized and well managed given constraints such as battery duration and real-time requirements for input–output operations.

The requirements for compute time of LL operations in category 2 depend on factors such as model size, input size, input–output frequency, algorithm-specific LL overheads and, crucially for ShELL, on how often distributions and tasks change, as well as how much knowledge can be acquired via sharing. As a result, suitable ShELL hardware platforms may require neural accelerators that can be dynamically allocated to various operations in category 2 (when learning), category 1 (during input–output operations), category 3 when extracting or integrating knowledge for/from sharing, or a trade-off of all these when performing inference, LL and communication at the same time. Variable bandwidth availability also demands dynamic allocation of compute power: at times when communication is not possible, more reliance on learning at the edge may be required. Therefore, category 4 may be also critical to dynamically partition the hardware and allocate priorities, for example, accelerating learning during periods of input–output inactivity such as maintenance or battery recharge, or during times when communication is not available.

In summary, a ShELL device requires a standard computer architecture with a particular focus on neural accelerators and parallel execution of operations with variable demands over time. Research and investments on new, fast, and efficient AI accelerators[128,142], particularly for lifelong learning[127] will facilitate deployment on the edge. A key performance factor in ShELL is determined by the ability of an edge device to optimize the timing and dynamic allocations of all operations illustrated in Fig. 4.

Emerging approaches to lifelong learning and sharing

Approaches that aim to implement the objectives in 'Synergy between lifelong learning and sharing' using the methods in 'Constituent fields of a lifelong learning collective' surveyed above have been proposed in recent years, demonstrating the potential of the synergy of lifelong learning and sharing. In one of the first studies[143], a collective LL algorithm (COLLA) is a network of agents that share knowledge in a distributed and decentralized manner. Each agent learns a local dictionary that reflects local knowledge based on the Alternating Direction Method of Multipliers (ADMM) algorithm[144], and shares that knowledge to improve learning on new tasks. A following study[145] improves upon the COLLA approach to allow each agent to maintain local agent-specific skills in addition to sharing collective knowledge.

Parameter isolation approaches based on modulating masks[40,146,147] have shown promise to facilitate task-specific transfer of knowledge. The idea is to exploit task-specific representations, encoded as network sub-regions (masks), to transfer specific machine-learned knowledge across agents. In ref. 94, the authors illustrate a lifelong reinforcement learning decentralized collective of agents, in which agents query each other for modulating masks and transfer those that are relevant to the current task on a peer-to-peer basis. One limitation is that agents require the same network backbone to be able to share masks, but the fixed backbone allows for knowledge reuse via a linear combination of masks. A similar concept is also exploited in other studies[102,103,148–150], where task-agnostic backbones are combined with small sets of task-specific parameters that allow for knowledge transfer with minimal changes to a model[151,152].

Recent extensions to federated learning have resulted in ShELL-likes studies, although often still based on centralized approaches. The methods in ref. 153 and ref. 154 proposed consecutive learning of classic ML datasets with an LL component in federated learning, allowing for more effective learning of non-IID datasets. Another extension of federated learning[57,155] decomposes the network weights into global federated parameters and sparse task-specific parameters, which allow agent-specific and personalized tasks. In ref. 60, federated learning is extended to account for both data and model heterogeneity.

Distillation methods have been used in a continual federated learning approach to diffuse computation across edge devices[156]. Along with knowledge transfer, distillation can also be utilized to reduce the size of models for use on edge devices, while still retaining equivalent performance. Due to these favourable qualities, knowledge distillation has been used in different approaches to implementing LL and sharing. Distillation approaches[157,158] may involve data sharing (Fig. 3) that potentially violates privacy. However, recent studies have shown that synthetic and privacy-preserving data can be used[159,160]. Accompanying policies with representative data (that is, core sets, special cases and so on) could lead to a more effective knowledge transfer, particularly for heterogeneous systems[156,161]. Exploring selective experience replay[93] via shareable experience replay buffers, recent studies[162–165] demonstrate a set of methods to assist radiologists in localizing anatomy landmarks. The issue of scalability to large networks was addressed with the proposal of a hub-based system in which some nodes act as hubs to aggregate knowledge from local sub-networks.

The advantage of a collective of robots that share risks to maintain safety has been explored previously[166,167]. Exploiting the concept of shield in safe reinforcement learning[168], a dataset is created online as agents experience catastrophic actions, and share that with their peers. The obvious benefit is that costly mistakes need not be repeated if experience sharing is used. Moreover, coordination of exploration policies could lead to more efficient, faster and safer learning[169]. Further technical details of the studies cited in this section are provided in the Supplementary Information.

Application areas for ShELL technology

Application areas that can benefit from ShELL methods can be identified by four features: (1) problems that have distributed and sequential data, and are thus likely to observe changing distributions across locations and time; (2) a fast learning response in front of changes in the problem or environment is required; (3) local knowledge or policies are either desired or required; (4) execution on SWaP-constrained devices may be critical in some operation scenarios, for example, in remote locations or with reduced communication. Figure 5 lists application areas, task categories and requirements that can benefit from ShELL. Task feature (4) is present only for some applications because learning and sharing can be beneficial regardless of hardware and domain constraints. Four exemplary domains are further described in the following sections.

Multi-agent active sensing

Multi-agent active sensing (MAAS) refers to the problem of coordinating multiple agents to strategically gather information from an underlying domain of interest to collectively achieve a desired sensing objective, for example, target detection or generative modelling of the domain[170,171]. MAAS is critical in many applications, from search and rescue to localization and anomaly detection in military reconnaissance problems[172]. With the recent advances in implicit neural representations (INRs)[173], the neural radiance fields (NeRFs)[174,175] and signed distance functions (SDFs)[176,177], there is an emerging interest in using such INRs in MAAS. We note that ShELL technology could enable a swarm of agents to collectively and actively construct an INR for the domain of interest while being resilient to adversarial and environmental disruptions, as demonstrated previously[178]. ShELL is particularly useful in this setting owing to (1) the need for decentralized coordination and sharing for optimal model construction, (2) limited communication bandwidth and (3) SWaP constraints that emerge from requiring a large, yet dispensable, swarm of agents to cover vast areas of interest while maintaining redundancy for resilience against adversities.

Space exploration

Space exploration has been identified as a domain in which ML can provide significant advantages[179]. While an overview of all current ML-aided tasks in space exploration is outside the scope of this paper, we note that ShELL technology can be beneficial in this domain due to (1) large delays and limited bandwidth between distant locations[180]; (2) the requirement for autonomous navigation systems and spacecraft control[181], ML-aided sensing[182,183] and autonomous decision-making; (3) the need for self-adaptation due to faults, unforeseen or changing conditions; (4) the limited power and computation available on spacecraft. For example, with a round-trip light-time of 6 to 40 minutes[184], communication to Mars and beyond is severely limited. Satellites and crafts powered by solar panels are also limited in their computational resources, but the need to adapt to unforeseen conditions is essential when human intervention is significantly reduced, delayed or absent. Thus, the ability to lifelong learn and effectively communicate with a fleet of orbiters and rovers will be a key to a successful Mars exploration effort.

Responsive and personalized medicine

The application of ML to medical domains has recently benefited from distributed approaches based on federated learning[185], in particular through enabling multi-institution collaborations[186]. However, a patient-centred approach, rather than tackling a single pathology, might require knowledge of a large variety of possible conditions and data acquisition domains (for example, genomics and imaging). If such knowledge is available in medical datasets, ShELL technology could be the key to delivering personalized, task-specific knowledge to the point of care[187]. Moreover, a constant evolution of illnesses and pathogens and new advanced diagnostic technologies limit large, centralized and single-task models to be responsive to new conditions. In such a context, ShELL agents that can asynchronously share experiences with each other to collectively learn multiple tasks over time offer a significant advantage over centralized single-task models[164]. A previous study[163] demonstrates the use of deep reinforcement learning on edge devices to continually learn and adapt to changing medical imaging environments and adjust to low-compute devices.

Responsive distributed cyber-security systems

Concerns over the impact of cyber attacks and cyber warfare continuously increase as more and more aspects of our lives depend on connected devices[188]. ML applications to cyber security have grown in recent years with scientific efforts and integration into commercial products[189–191]. One challenge in this domain is that threats evolve continuously, and vary across different types of devices and locations. Effective responses are required at a fast time scale. This makes it difficult to train ML systems on static datasets, making such scenarios suitable for lifelong learning[192]. The distributed nature of network devices also implies that information on malicious activity is first recorded at the edge: data gathering for learning centralized models might not be effective. The deployment of ShELL-like systems enables continuous adaptation to evolving threats and, crucially, fast communication of

Application area	Task category	Learning tasks (LL and sharing)
Common features		**Distributed and sequential data** with changing distributions
		Fast response/learning time desired/required
		Local policies desired/required
		Learning on **edge devices** in remote locations or with reduced communication
Robotics	Search and rescue	New environments in disaster-response search and resuce operations and fast learning requirements
	Disaster response	New unforeseen objectives, different environments and requirement for rapid on-site learning
	Surveillance	New monitoring targets and features, evolving distributions of inputs with distributed sensing
	Human–robot interaction	Personalized assistance tasks: changing user-dependent distributions and tasks
		New or unforeseen collaborative and personalized tasks
	Manufacturing	New components, tools or assembly procedures to learn and use across multiple robots
		Unforeseen outliers in multi-robot on-site inspection operations
		New observations in quality control due to new products or requirements
		New maintenance task with unknown policies to be rapidly acquired by multi-robot systems
	Autonomous vehicles	Unforeseen conditions requiring new local policies that can be shared with the fleet
	Space exploration	Reduced, delayed, or costly communication requiring edge learning with occasional sharing
		Small edge devices with reduced power, unknown conditions, possible faults and errors
Cyber security	Defence	New and evolving threats for deployed multi-agent autonomous defence systems
	Crime and fraud	Fast evolving scenarios, features and data across locations or on distributed infrastructures
		Learning communication patterns for cyber-security in IoT devices
		New unforeseen anomalies and patterns in networks revealing malicious activity
Digital agents	ML agents	New personalized tasks for personal assistants integrating local and global knowledge
		Fast evolving domain-specific and local information (e.g. finance, law) across locations
	Language personalization	Learning new evolving local language patterns and expressions
IoT	Adaptive IoT devices	Learning and sharing exceptions across large numbers of devices
Medicine	Diagnosis	New patient data for emerging health threats with local data and distributions
	Personalized medicine	Integration of patient-specific and site-specific information with global knowledge
	Triage operations	New sudden emergency with distributed data collection requiring fast response time

Fig. 5 | List of application areas, task categories and learning tasks that are suitable to ShELL systems. The first three features are common to all identified areas because they summarize the typical ShELL-suitable domains, that is, distributed and sequential data with changing distributions, fast response (learning) time desired/required, and local or personalized policies required/desired. The last common feature, learning on edge devices in remote locations with reduced communication, may be found in some application areas.

defence strategies across the collective of distributed agents. One key advantage of ShELL-like systems is the ability to observe a new threat or vulnerability at one particular node, discover an appropriate defence policy, and communicate that to all other nodes in the collective, which then becomes rapidly immune to that particular attack.

Outstanding challenges and opportunities

While studies described in 'Emerging approaches to lifelong learning and sharing' illustrate promising initial implementations of ShELL methods, further developments of large ShELL-like systems present open challenges.

- Scalability. Managing the connectivity of a potentially large number of agents is essential to achieve scalability[193,194]. Just as information on the Internet is searched and transferred with the help of search engines and large data distribution centres, similarly, a collective of ML agents, while decentralized, might require hubs or communication nodes[162,195]. A perhaps even more critical aspect of scalability is the organization of knowledge and how to answer questions such as which agents know what: algorithms that define, label and organize tasks[86] or search the collective for specific knowledge might be required.
- Protocols. The methods reviewed in 'Emerging approaches to lifelong learning and sharing' reveal that the choice of what data is transferred across agents, and their timing, can vary across different ShELL algorithms, implying that a unified approach has not yet been identified, or may not be desirable given different requirements and domains. The different types of data transfer highlighted in Fig. 3 suggest that standards, protocols and possibly even a language, are required to enable a widespread distribution of machine-learned knowledge.
- Computation. The development of AI accelerators in hardware will be key to determining the rate of diffusion of learning at the edge[127]. Current trends suggest that learning at the edge will become more feasible, although computing power alone will not be sufficient to solve the algorithmic challenges of ShELL that will ultimately determine its effectiveness and widespread use.

Despite the aforementioned challenges, ShELL algorithms have the potential to augment the capabilities of the latest and most advanced AI models, including transformer networks[196,197] and foundation models[198,199]. The increasingly large amount of data and the computing power required to train foundation models may limit their scalability, and already limit the number of entities that possess sufficient resources for full re-training (for example, development and training cost of GPT-4 exceeded US$100 million[200]). If such models can integrate new knowledge over a lifetime, and request knowledge from other models, the need to re-train and use large amounts of data could be reduced. When required, knowledge from similar models could be retrieved on particular topics for example, the specific laws of a country, the policies of one organization, or new scientific or technical knowledge recently developed.

A collective of ShELL agents has the potential to give rise to a worldwide network of AI. Therefore, it is essential to take a long-term view of the risks[201] of such a possibility. We highlight the following two main risks and suggest related mitigating approaches.

- Malicious attacks or errors. A learning and sharing collective may be vulnerable to malicious attacks: knowledge obtained from other agents might be incorrect or set up to deceive. While this is not a new problem when dealing with information, the fast and automatic nature of a ShELL collective could lead to worse consequences. Solutions could involve relying on a close collective of trusted agents. Even more effectively, knowledge acquired from other agents could be tested and validated before use. As ML applications are increasingly adopted in different areas of society, it may be beneficial for agents to incorporate mechanisms to verify the alignment of acquired knowledge with their objectives. This will be essential in mission-critical operations where safety and compliance with specific standards need to be ensured. A governance issue[202] might also emerge if each agent in the population can obtain any information, expand it, and use it for individual objectives. As a consequence, since a decentralized collective of agents has no central governance, responsibility, and accountability might be unclear.

- Fast and autonomous knowledge sharing. The potential for an AI agent to acquire and disseminate unethical or illicit capabilities at a very fast time scale, before human intervention, poses a significant concern. Eliminating knowledge from a number of agents or terminating them may not be sufficient to remove or control knowledge that has spread across a collective. A similar issue is already present with the Internet, which makes eradicating specific information particularly difficult once it has spread. In addition, the processes used by LL agents when integrating new knowledge in their models may be difficult to reverse, that is, to unlearn[203]. Although the alignment problem[204] is not specific to ShELL, the rapidity of spreading malicious information or skills across a ShELL collective could amplify the problem. Entities that integrate knowledge over a lifetime, continuously adapt to new situations, and rapidly share new knowledge with all their peers, like in ShELL, have been proposed in science fiction to voice concerns over collective AI-related risks (for example, the Borg collective from the Star Trek franchise). However, we posit that a key safety aspect in the research highlighted in this paper is a decentralized structure of independent and autonomous agents with possibly different goals. Sharing among LL agents that maintain individual objectives and agency could provide an efficient, diverse and resilient democracy of AI to contrast the emergence of few, large and centralized AI models[205,206]. Efforts to categorize and label specific knowledge for optimized sharing may improve transparency and result in a more explainable structure of machine-learned knowledge or policies, for example, through modular or component-wise isolation[98].

Conclusion

The ideas and research in this paper highlight new trends that point to a better reuse of machine-learned knowledge. The reuse can happen within a lifelong learning agent that learns incrementally without forgetting[98], and among agents that exchange machine-learned knowledge. Converging fields and recent studies surveyed in this Perspective demonstrate the feasibility of such concepts, but a widespread use of the technology will depend on the convergence to common AI communication protocols for LL and further developments of hardware to learn at the edge. Knowledge reuse could lead to long-term scalability of AI that may be fully achieved only if machine-learned knowledge can be added incrementally and shared among agents, resulting also in the reduction of energy and carbon footprint of increasingly large systems[25,26,200,207]. The autonomous reuse and sharing of machine-learned knowledge has the potential to significantly accelerate progress in AI capabilities and see the emergence of more powerful systems that retrieve, integrate and build upon existing knowledge to achieve more complex goals at significantly shorter time scales.

References

1. Fagan, M. Collective scientific knowledge. *Philos. Compass* **7**, 821–831 (2012).
2. Csibra, G. & Gergely, G. Natural pedagogy as evolutionary adaptation. *Phil. Trans. R. Soc. B* **366**, 1149–1157 (2011).
3. Wooldridge, M. & Jennings, N. R. Intelligent agents: theory and practice. *Knowl. Eng. Rev.* **10**, 115–152 (1995).
4. Ferber, J. & Weiss, G. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence* Vol. 1 (Addison-Wesley, 1999).
5. Stone, P. & Veloso, M. Multiagent systems: a survey from a machine learning perspective. *Auton. Rob.* **8**, 345–383 (2000).
6. Conitzer, V. & Oesterheld, C. Foundations of cooperative AI. In *Proc. AAAI Conference on Artificial Intelligence* Vol. 37, 15359–15367 (AAAI, 2022).
7. Semsar-Kazerooni, E. & Khorasani, K. Multi-agent team cooperation: a game theory approach. *Automatica* **45**, 2205–2213 (2009).
8. Thrun, S. Is learning the *n*-th thing any easier than learning the first? In *Advances in Neural Information Processing Systems* Vol. 8 (1995).
9. Thrun, S. Lifelong learning algorithms. *Learning to Learn* **8**, 181–209 (1998).
10. Chen, Z. & Liu, B. *Lifelong Machine Learning* Vol. 1 (Springer, 2018).
11. Kudithipudi, D. et al. Biological underpinnings for lifelong learning machines. *Nat. Mach. Intell.* **4**, 196–210 (2022).
12. Mundt, M., Hong, Y., Pliushch, I. & Ramesh, V. A wholistic view of continual learning with deep neural networks: forgotten lessons and the bridge to active and open world learning. *Neural Networks* **160**, 306–336 (2023).
13. Khetarpal, K., Riemer, M., Rish, I. & Precup, D. Towards continual reinforcement learning: a review and perspectives. *J. Artif. Intell. Res.* **75**, 1401–1476 (2022).
14. Mendez, J. A., van Seijen, H. & Eaton, E. Modular lifelong reinforcement learning via neural composition. In *International Conference on Learning Representations* (2022).
15. Li, T., Sahu, A. K., Talwalkar, A. & Smith, V. Federated learning: challenges, methods, and future directions. *IEEE Signal Process. Mag.* **37**, 50–60 (2020).
16. Dorri, A., Kanhere, S. S. & Jurdak, R. Multi-agent systems: a survey. *IEEE Access* **6**, 28573–28593 (2018).
17. Shi, W., Cao, J., Zhang, Q., Li, Y. & Xu, L. Edge computing: vision and challenges. *IEEE Internet Things J.* **3**, 637–646 (2016).
18. Cai, H. et al. Enable deep learning on mobile devices: methods, systems, and applications. *ACM Trans. Des. Autom. Electron. Syst.* **27**, 20 (2022).
19. Shared-Experience Lifelong Learning (ShELL). Opportunity DARPA-PA-20-02-11. *SAM.gov* https://sam.gov/opp/1afbf600f2e04b26941fad352c08d1f1/view (accessed 10 October 2023).
20. Smith, P. et al. Network resilience: a systematic approach. *IEEE Commun. Mag.* **49**, 88–97 (2011).
21. Zhang, J., Cheung, B., Finn, C., Levine, S. & Jayaraman, D. Cautious adaptation for reinforcement learning in safety-critical settings. In *International Conference on Machine Learning* 11055–11065 (PMLR, 2020).
22. McMahan, B., Moore, E., Ramage, D., Hampson, S. & Arcas, B. A. Y. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics* 1273–1282, (PMLR, 2017).
23. Liu, J. et al. From distributed machine learning to federated learning: a survey. *Knowl. Inf. Syst.* **64**, 885–917 (2022).
24. Verbraeken, J. et al. A survey on distributed machine learning. *ACM Comput. Surv.* **53**, 30 (2020).
25. Henderson, P. et al. Towards the systematic reporting of the energy and carbon footprints of machine learning. *J. Mach. Learn. Res.* **21**, 10039–10081 (2020).
26. de Vries, A. The growing energy footprint of artificial intelligence. *Joule* **7**, 2191–2194 (2023).
27. Silver, D. L., Yang, Q. & Li, L. Lifelong machine learning systems: beyond learning algorithms. In *2013 AAAI Spring Symposium Series* (AAAI, 2013).

28. Hadsell, R., Rao, D., Rusu, A. A. & Pascanu, R. Embracing change: continual learning in deep neural networks. *Trends Cognit. Sci.* **24**, 1028–1040 (2020).
29. French, R. M. Catastrophic forgetting in connectionist networks. *Trends Cognit. Sci.* **3**, 128–135 (1999).
30. Kirkpatrick, J. et al. Overcoming catastrophic forgetting in neural networks. *Proc. Natl Acad. Sci. USA* **114**, 3521–3526 (2017).
31. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C. & Wermter, S. Continual lifelong learning with neural networks: a review. *Neural Networks* **113**, 54–71 (2018).
32. Lange, M. D. et al. A continual learning survey: defying forgetting in classification tasks. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 3366–3385 (2022).
33. van de Ven, G. M., Tuytelaars, T. & Tolias, A. S. Three types of incremental learning. *Nat. Mach. Intell.* **4**, 1185–1197 (2022).
34. Soltoggio, A., Stanley, K. O. & Risi, S. Born to learn: the inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks* **108**, 48–67 (2018).
35. Lifelong learning machines (L2M). *DARPA* https://www.darpa.mil/news-events/2017-03-16 (accessed 10 October 2023).
36. New, A., Baker, M., Nguyen, E. & Vallabha, G. Lifelong learning metrics. Preprint at https://doi.org/10.48550/arXiv.2201.08278 (2022).
37. Baker, M. M. et al. A domain-agnostic approach for characterization of lifelong learning systems. *Neural Networks* **160**, 274–296 (2023).
38. Mendez, J. A. & Eaton, E. Lifelong learning of compositional structures. In *International Conference on Learning Representations* (2021).
39. Xie, A. & Finn, C. Lifelong robotic reinforcement learning by retaining experiences. In *Conference on Lifelong Learning Agents* 838–855 (PMLR, 2022).
40. Ben-Iwhiwhu, E., Nath, S., Pilly, P. K., Kolouri, S. & Soltoggio, A. Lifelong reinforcement learning with modulating masks. In *Transactions on Machine Learning Research* (2023).
41. Tasse, G. N., James, S. & Rosman, B. Generalisation in lifelong reinforcement learning through logical composition. In *International Conference on Learning Representations* (2022).
42. Merenda, M., Porcaro, C. & Iero, D. Edge machine learning for AI-enabled IoT devices: a review. *Sensors* **20**, 2533 (2020).
43. Sipola, T., Alatalo, J., Kokkonen, T. & Rantonen, M. Artificial intelligence in the IoT era: a review of edge AI hardware and software. In *2022 31st Conference of Open Innovations Association (FRUCT)* 320–331 (IEEE, 2022).
44. Prabhu, A. et al. Computationally budgeted continual learning: What does matter? In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 3698–3707 (2023).
45. Díaz-Rodríguez, N., Lomonaco, V., Filliat, D. & Maltoni, D. Don't forget, there is more than forgetting: new metrics for continual learning. Preprint at https://doi.org/10.48550/arXiv.1810.13166 (2018).
46. De Lange, M., van de Ven, G. & Tuytelaars, T. Continual evaluation for lifelong learning: identifying the stability gap. In *11th International Conference on Learning Representations* https://openreview.net/forum?id=Zy350cRstc6 (ICLR, 2023).
47. Ghunaim, Y. et al. Real-time evaluation in online continual learning: a new hope. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 11888–11897 (2023).
48. Sarker, I. H. Machine learning: algorithms, real-world applications and research directions. *SN Comput. Sci.* **2**, 160 (2021).
49. Tsuda, B., Tye, K. M., Siegelmann, H. T. & Sejnowski, T. J. A modeling framework for adaptive lifelong learning with transfer and savings through gating in the prefrontal cortex. *Proc. Natl Acad. Sci. USA* **117**, 29872–29882 (2020).
50. Kairouz, P. et al. Advances and open problems in federated learning. *Found. Trends Mach. Learn.* **14**, 1–210 (2021).
51. Zhu, H., Xu, J., Liu, S. & Jin, Y. Federated learning on non-IID data: a survey. *Neuropcomputing* **465**, 371–390 (2021).
52. Nguyen, D. C. et al. Federated learning for internet of things: a comprehensive survey. *IEEE Commun. Surv. Tutorials* **23**, 1622–1658 (2021).
53. Abreha, H. G., Hayajneh, M. & Serhani, M. A. Federated learning in edge computing: a systematic survey. *Sensors* **22**, 450 (2022).
54. Guo, Y., Lin, T. & Tang, X. Towards federated learning on time-evolving heterogeneous data. Preprint at https://doi.org/10.48550/arXiv.2112.13246 (2021).
55. Criado, M. F., Casado, F. E., Iglesias, R., Regueiro, C. V. & Barro, S. Non-IID data and continual learning processes in federated learning: a long road ahead. *Inf. Fusion* **88**, 263–280 (2022).
56. Gao, L. et al. FedDC: federated learning with non-IID data via local drift decoupling and correction. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 10112–10121 (2022).
57. Yoon, J., Jeong, W., Lee, G., Yang, E. & Hwang, S. J. Federated continual learning with weighted inter-client transfer. In *International Conference on Machine Learning* 12073–12086 (PMLR, 2021).
58. Pellegrini, L., Lomonaco, V., Graffieti, G. & Maltoni, D. Continual learning at the edge: real-time training on smartphone devices. In *Proc. European Symposium on Artificial Neural Networks* https://doi.org/10.14428/esann/2021.ES2021-136 (2021).
59. Gao, D. et al. Rethinking pruning for accelerating deep inference at the edge. In *Proc. 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 155–164 (2020).
60. Huang, W., Ye, M. & Du, B. Learn from others and be yourself in heterogeneous federated learning. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 10143–10153 (2022).
61. Sun, T., Li, D. & Wang, B. Decentralized federated averaging. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 4289–4301 (2023).
62. Taylor, M. E. & Stone, P. Transfer learning for reinforcement learning domains: a survey. *J. Mach. Learn. Res.* **10**, 1633–1685 (2009).
63. Zamir, A. R. et al. Taskonomy: disentangling task transfer learning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (2018).
64. Zhuang, F. et al. A comprehensive survey on transfer learning. *Proc. IEEE* **109**, 43–76 (2020).
65. Ding, N. et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nat. Mach. Intell.* **5**, 220–235 (2023).
66. Koohpayegani, S. A., Navaneet, K., Nooralinejad, P., Kolouri, S. & Pirsiavash, H. NOLA: networks as linear combination of low rank random basis. In *International Conference on Learning Representations* (ICLR, 2024).
67. Wang, M. & Deng, W. Deep visual domain adaptation: a survey. *Neurocomputing* **312**, 135–153 (2018).
68. Wilson, G. & Cook, D. J. A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.* **11**, 51 (2020).
69. Farahani, A., Voghoei, S., Rasheed, K. & Arabnia, H. R. A brief review of domain adaptation. In *Advances in Data Science and Information Engineering: Proceedings from ICDATA 2020 and IKE 2020* 877–894 (2021).
70. Kim, Y., Cho, D., Han, K., Panda, P. & Hong, S. Domain adaptation without source data. *IEEE Trans. Artif. Intell.* **2**, 508–518 (2021).
71. Caruana, R. Multitask learning. *Mach. Learn.* **28**, 41–75 (1997).
72. Luong, M.-T., Le, Q. V., Sutskever, I., Vinyals, O. & Kaiser, L. Multi-task sequence to sequence learning. Preprint at https://doi.org/10.48550/arXiv.1511.06114 (2015).
73. Ruder, S. An overview of multi-task learning in deep neural networks. Preprint at https://doi.org/10.48550/arXiv.1706.05098 (2017).

74. Liu, X., He, P., Chen, W. & Gao, J. Multi-task deep neural networks for natural language understanding. In *Proc. 57th Annual Meeting of the Association for Computational Linguistics* (2019).
75. Hospedales, T., Antoniou, A., Micaelli, P. & Storkey, A. Meta-learning in neural networks: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 5149–5169 (2021).
76. Kayaalp, M., Vlaski, S. & Sayed, A. H. Dif-MAML: decentralized multi-agent meta-learning. *IEEE Open J. Signal Process.* **3**, 71–93 (2022).
77. Riemer, M. et al. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *7th International Conference on Learning Representations, ICLR 2019* (OpenReview, 2019).
78. Bengio, Y., Louradour, J., Collobert, R. & Weston, J. Curriculum learning. In *Proc. 26th Annual International Conference on Machine Learning* 41–48 (2009).
79. Narvekar, S. et al. Curriculum learning for reinforcement learning domains: a framework and survey. *J. Mach. Learn. Res.* **21**, 7382–7431 (2020).
80. Wang, W., Zheng, V. W., Yu, H. & Miao, C. A survey of zero-shot learning: settings, methods, and applications. *ACM Trans. Intell. Syst. Technol.* **10**, 13 (2019).
81. Rostami, M., Isele, D. & Eaton, E. Using task descriptions in lifelong machine learning for improved performance and zero-shot transfer. *J. Artif. Intell. Res.* **67**, 673–704 (2020).
82. Chen, J. et al. Knowledge-aware zero-shot learning: survey and perspective. In *Proc. 30th International Joint Conference on Artificial Intelligence (IJCAI-21)* (2021).
83. Xie, G.-S., Zhang, Z., Xiong, H., Shao, L. & Li, X. Towards zero-shot learning: a brief review and an attention-based embedding network. *IEEE Trans. Circuits Syst. Video Technol.* **33**, 1181–1197 (2022).
84. Cao, W. et al. A review on multimodal zero-shot learning. *Wiley Interdiscip. Rev. Data Min. Knowl. Discovery* **13**, e1488 (2023).
85. Jones, A. M. et al. USC-DCT: a collection of diverse classification tasks. *Data* **8**, 153 (2023).
86. Liu, X., Bai, Y., Lu, Y., Soltoggio, A. & Kolouri, S. Wasserstein task embedding for measuring task similarities. Preprint at https://doi.org/10.48550/arXiv.2208.11726 (2022).
87. Yang, J., Zhou, K., Li, Y. & Liu, Z. Generalized out-of-distribution detection: a survey. Preprint at https://doi.org/10.48550/arXiv.2110.11334 (2021).
88. Abdar, M. et al. A review of uncertainty quantification in deep learning: techniques, applications and challenges. *Inf. Fusion* **76**, 243–297 (2021).
89. Musliner, D. J. et al. OpenMIND: planning and adapting in domains with novelty. In *Proc. 9th Conference on Advances in Cognitive Systems* (2021).
90. Rios, A. & Itti, L. Lifelong learning without a task oracle. In *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence* 255–263 (IEEE, 2020).
91. Carvalho, D. V., Pereira, E. M. & Cardoso, J. S. Machine learning interpretability: a survey on methods and metrics. *Electronics* **8**, 832 (2019).
92. Masana, M. et al. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 5513–5533 (2022).
93. Isele, D. & Cosgun, A. Selective experience replay for lifelong learning. In *Proc. AAAI Conference on Artificial Intelligence* Vol. 32 (2018).
94. Nath, S. et al. Sharing lifelong reinforcement learning knowledge via modulating masks. In *Proc. of Machine Learning Research* Vol. 232 (2023).
95. Pimentel, M. A., Clifton, D. A., Clifton, L. & Tarassenko, L. A review of novelty detection. *Signal Process.* **99**, 215–249 (2014).
96. Da Silva, B. C., Basso, E. W., Bazzan, A. L. & Engel, P. M. Dealing with non-stationary environments using context detection. In *Proc. 23rd International Conference on Machine Learning* 217–224 (2006).
97. Niv, Y. Learning task-state representations. *Nat. Neurosci.* **22**, 1544–1553 (2019).
98. Mendez, J. & Eaton, E. How to reuse and compose knowledge for a lifetime of tasks: a survey on continual learning and functional composition. In *Transactions on Machine Learning Research* (2023).
99. Hu, E. J. et al. LoRA: low-rank adaptation of large language models. *International Conference on Learning Representations (ICLR)* (2021).
100. Nooralinejad, P. et al. PRANC: pseudo random networks for compacting deep models. In *Proc. IEEE/CVF International Conference on Computer Vision* 17021–17031 (2023).
101. Lester, B., Al-Rfou, R. & Constant, N. The power of scale for parameter-efficient prompt tuning. In *Proc. 2021 Conference on Empirical Methods in Natural Language Processing* (2021).
102. Ge, Y. et al. Lightweight learner for shared knowledge lifelong learning. In *Transactions on Machine Learning Research* (2023).
103. Ge, Y. et al. CLR: Channel-wise lightweight reprogramming for continual learning. In *Proc. IEEE/CVF International Conference on Computer Vision* 18798–18808 (2023).
104. Sarker, M. K., Zhou, L., Eberhart, A. & Hitzler, P. Neuro-symbolic artificial intelligence. *AI Commun.* **34**, 197–209 (2021).
105. Zoph, B. & Le, Q. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations* (2017).
106. Ren, P. et al. A comprehensive survey of neural architecture search: challenges and solutions. *ACM Comput. Surv.* **54**, 76 (2021).
107. Zhang, C., Patras, P. & Haddadi, H. Deep learning in mobile and wireless networking: a survey. *IEEE Commun. Surv. Tutorials* **21**, 2224–2287 (2019).
108. Deng, S. et al. Edge intelligence: the confluence of edge computing and artificial intelligence. *IEEE Internet Things J.* **7**, 7457–7469 (2020).
109. Murshed, M. S. et al. Machine learning at the network edge: a survey. *ACM Comput. Surv.* **54**, 170 (2021).
110. Ajani, T. S., Imoize, A. L. & Atayero, A. A. An overview of machine learning within embedded and mobile devices–optimizations and applications. *Sensors* **21**, 4412 (2021).
111. Dhar, S. et al. A survey of on-device machine learning: an algorithms and learning theory perspective. *ACM Trans. Internet Things* **2**, 15 (2021).
112. Singh, R. & Gill, S. S. Edge AI: a survey. *Internet Things Cyber-Phys. Syst.* **3**, 71–92 (2023).
113. Mao, Y., You, C., Zhang, J., Huang, K. & Letaief, K. B. A survey on mobile edge computing: the communication perspective. *IEEE Commun. Surv. Tutorials* **19**, 2322–2358 (2017).
114. Xu, D. et al. Edge intelligence: architectures, challenges, and applications. Preprint at https://doi.org/10.48550/arXiv.2003.12172 (2020).
115. Li, E., Zeng, L., Zhou, Z. & Chen, X. Edge AI: on-demand accelerating deep neural network inference via edge computing. *IEEE Trans. Wireless Commun.* **19**, 447–457 (2019).
116. Mehlin, V., Schacht, S. & Lanquillon, C. Towards energy-efficient deep learning: an overview of energy-efficient approaches along the deep learning lifecycle. Preprint at https://doi.org/10.48550/arXiv.2303.01980 (2023).
117. Lin, J. et al. On-device training under 256kb memory. 36th Conference on Neural Information Processing Systems (NeurIPS) (2022).
118. Yang, Y., Li, G. & Marculescu, R. Efficient on-device training via gradient filtering. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 3811–3820 (2023).

119. Hayes, T. L. & Kanan, C. Online continual learning for embedded devices. In *Proc. First Conference on Lifelong Learning Agents* (eds Chandar, S. et al.) 744–766 (PMLR, 2022).
120. Wang, Z. et al. SparCL: sparse continual learning on the edge. In *36th Conference on Neural Information Processing Systems* (2022).
121. Harun, M. Y., Gallardo, J., Hayes, T. L. & Kanan, C. How efficient are today's continual learning algorithms? In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2430–2435 (2023).
122. Yang, J. et al. Quantization networks. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 7308–7316 (2019).
123. Cai, Z., He, X., Sun, J. & Vasconcelos, N. Deep learning with low precision by half-wave gaussian quantization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* 5918–5926 (2017).
124. Jain, A., Bhattacharya, S., Masuda, M., Sharma, V. & Wang, Y. Efficient execution of quantized deep learning models: a compiler approach. Preprint at https://doi.org/10.48550/arXiv.2006.10226 (2020).
125. Goel, A., Tung, C., Lu, Y.-H. & Thiruvathukal, G. K. A survey of methods for low-power deep learning and computer vision. In *2020 IEEE 6th World Forum on Internet of Things* (IEEE, 2020).
126. Ma, X. et al. Cost-effective on-device continual learning over memory hierarchy with Miro. In *Proc. 29th Annual International Conference on Mobile Computing and Networking* 83, 1–15 (ACM, 2023).
127. Kudithipudi, D. et al. Design principles for lifelong learning AI accelerators. *Nat. Electron.* **6**, 807–822 (2023).
128. Machupalli, R., Hossain, M. & Mandal, M. Review of ASIC accelerators for deep neural network. *Microprocess. Microsyst.* **89**, 104441 (2022).
129. Jouppi, P. N. et al. In-datacenter performance analysis of a tensor processing unit. In *Proc. 44th Annual International Symposium on Computer Architecture* (2017).
130. Sebastian, A., Le Gallo, M., Khaddam-Aljameh, R. & Eleftheriou, E. Memory devices and applications for in-memory computing. *Nat. Nanotechnol.* **15**, 529–544 (2020).
131. Tang, K.-T. et al. Considerations of integrating computing-in-memory and processing-in-sensor into convolutional neural network accelerators for low-power edge devices. In *2019 Symposium on VLSI Circuits* T166–T167 (IEEE, 2019).
132. Roy, K., Jaiswal, A. & Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature* **575**, 607–617 (2019).
133. Chakraborty, I., Jaiswal, A., Saha, A., Gupta, S. & Roy, K. Pathways to efficient neuromorphic computing with non-volatile memory technologies. *Appl. Phys. Rev.* **7**, 021308 (2020).
134. Christensen, D. V. et al. 2022 roadmap on neuromorphic computing and engineering. *Neuromorph. Comput. Eng.* **2**, 022501 (2022).
135. Rathi, N. et al. Exploring neuromorphic computing based on spiking neural networks: algorithms to hardware. *ACM Comput. Surv.* **55**, 243 (2023).
136. Zhang, W. et al. Neuro-inspired computing chips. *Nat. Electron.* **3**, 371–382 (2020).
137. Feldmann, J. et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021).
138. Peserico, N., Shastri, B. J. & Sorger, V. J. Integrated photonic tensor processing unit for a matrix multiply: a review. *J. Lightwave Technol.* **41**, 3704–3716 (2023).
139. Shastri, B. J. et al. Photonics for artificial intelligence and neuromorphic computing. *Nat. Photonics* **15**, 102–114 (2021).
140. Toczé, K. & Nadjm-Tehrani, S. A taxonomy for management and optimization of multiple resources in edge computing. *Wireless Commun. Mobile Comput.* **2018**, 7476201 (2018).
141. Bhattacharjee, A., Venkatesha, Y., Moitra, A. & Panda, P. MIME: adapting a single neural network for multi-task inference with memory-efficient dynamic pruning. In *Proc. 59th ACM/IEEE Design Automation Conference* 499–504 (2022).
142. Extreme Computing BAA. *DARPA* https://sam.gov/opp/211b1819bd5f46eba20d4a466358d8bb/view (accessed 10 October 2023).
143. Rostami, M., Kolouri, S., Kim, K. & Eaton, E. Multi-agent distributed lifelong learning for collective knowledge acquisition. In *Proc. 17th International Conference on Autonomous Agents and Multiagent Systems 2018* (2017).
144. Boyd, S. et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**, 1–122 (2011).
145. Mohammadi, J. & Kolouri, S. Collaborative learning through shared collective knowledge and local expertise. In *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing* (2019).
146. Wortsman, M. et al. Supermasks in superposition. *Adv. Neural Inf. Process. Syst.* **33**, 15173–15184 (2020).
147. Koster, N., Grothe, O. & Rettinger, A. Signing the supermask: keep, hide, invert. In *International Conference on Learning Representations* https://openreview.net/forum?id=e0jtGTfPihs (2022).
148. Wen, S., Rios, A., Ge, Y. & Itti, L. Beneficial perturbation network for designing general adaptive artificial intelligence systems. *IEEE Trans. Neural Networks Learn. Syst.* **33**, 3778–3791 (2021).
149. Saha, G., Garg, I. & Roy, K. Gradient projection memory for continual learning. In *International Conference on Learning Representations* (2021).
150. Choudhary, S., Aketi, S. A., Saha, G. & Roy, K. CoDeC: communication-efficient decentralized continual learning. Preprint at https://doi.org/10.48550/arXiv.2303.15378 (2023).
151. Singh, P., Verma, V. K., Mazumder, P., Carin, L. & Rai, P. Calibrating CNNs for lifelong learning. *Adv. Neural Inf. Process. Syst.* **33**, 15579–15590 (2020).
152. Verma, V. K., Liang, K. J., Mehta, N., Rai, P. & Carin, L. Efficient feature transformations for discriminative and generative continual learning. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 13865–13875 (2021).
153. Ma, Z., Lu, Y., Li, W. & Cui, S. EFL: elastic federated learning on non-IID data. In *Conference on Lifelong Learning Agents* 92–115 (PMLR, 2022).
154. Shenaj, D., Toldo, M., Rigon, A. & Zanuttigh, P. Asynchronous federated continual learning. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10208460 (2023).
155. Venkatesha, Y., Kim, Y., Park, H. & Panda, P. Divide-and-conquer the NAS puzzle in resource constrained federated learning systems. *Neural Networks* **168**, 569–579 (2023).
156. Usmanova, A., Portet, F., Lalanda, P. & Vega, G. A distillation-based approach integrating continual learning and federated learning for pervasive services. 3rd Workshop on Continual and Multimodal Learning for Internet of Things – Co-located with IJCAI 2021, Aug 2021, Montreal, Canada https://doi.org/10.48550/arXiv.2109.04197 (2021).
157. Wang, T., Zhu, J.-Y., Torralba, A. & Efros, A. A. Dataset distillation. Preprint at https://doi.org/10.48550/arXiv.1811.10959 (2018).
158. Cazenavette, G., Wang, T., Torralba, A., Efros, A. A. & Zhu, J.-Y. Dataset distillation by matching training trajectories. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 4750–4759 (2022).
159. Baradad Jurjo, M., Wulff, J., Wang, T., Isola, P. & Torralba, A. Learning to see by looking at noise. *Adv. Neural Inf. Process. Syst.* **34**, 2556–2569 (2021).

160. Carta, A., Cossu, A., Lomonaco, V., Bacciu, D. & van de Weijer, J. Projected latent distillation for data-agnostic consolidation in distributed continual learning. Preprint at https://doi.org/10.48550/arXiv.2303.15888 (2023).
161. Teh, Y. et al. Distral: robust multitask reinforcement learning. In *Advances in Neural Information Processing Systems* Vol. 30 (2017).
162. Zheng, G., Jacobs, M. A., Braverman, V. & Parekh, V. S. Asynchronous decentralized federated lifelong learning for landmark localization in medical imaging. In *International Workshop on Federated Learning for Distributed Data Mining* (2023).
163. Zheng, G., Lai, S., Braverman, V., Jacobs, M. A. & Parekh, V. S. A framework for dynamically training and adapting deep reinforcement learning models to different, low-compute, and continuously changing radiology deployment environments. Preprint at https://doi.org/10.48550/arXiv.2306.05310 (2023).
164. Zheng, G., Lai, S., Braverman, V., Jacobs, M. A. & Parekh, V. S. Multi-environment lifelong deep reinforcement learning for medical imaging. Preprint at https://doi.org/10.48550/arXiv.2306.00188 (2023).
165. Zheng, G., Zhou, S., Braverman, V., Jacobs, M. A. & Parekh, V. S. Selective experience replay compression using coresets for lifelong deep reinforcement learning in medical imaging. In *Proc. Machine Learning Research* 227, 1751–1764 (2024).
166. Shperberg, S. S., Liu, B. & Stone, P. Learning a shield from catastrophic action effects: never repeat the same mistake. Preprint at https://doi.org/10.48550/arXiv.2202.09516 (2022).
167. Shperberg, S. S., Liu, B., Allievi, A. & Stone, P. A rule-based shield: Accumulating safety rules from catastrophic action effects. In *Conference on Lifelong Learning Agents* 231–242 (PMLR, 2022).
168. Alshiekh, M. et al. Safe reinforcement learning via shielding. In *Proc. 32nd AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence* (AAAI, 2018).
169. Garcıa, J. & Fernández, F. A comprehensive survey on safe reinforcement learning. *J. Mach. Learn. Res.* **16**, 1437–1480 (2015).
170. Jang, D., Yoo, J., Son, C. Y., Kim, D. & Kim, H. J. Multi-robot active sensing and environmental model learning with distributed gaussian process. *IEEE Robot. Autom. Lett.* **5**, 5905–5912 (2020).
171. Igoe, C., Ghods, R. & Schneider, J. Multi-agent active search: a reinforcement learning approach. *IEEE Rob. Autom. Lett.* **7**, 754–761 (2021).
172. Raja, G., Baskar, Y., Dhanasekaran, P., Nawaz, R. & Yu, K. An efficient formation control mechanism for multi-UAV navigation in remote surveillance. In *2021 IEEE Globecom Workshops* (IEEE, 2021).
173. Sitzmann, V., Martel, J., Bergman, A., Lindell, D. & Wetzstein, G. Implicit neural representations with periodic activation functions. *Adv. Neural Inf. Process. Syst.* **33**, 7462–7473 (2020).
174. Yu, A., Ye, V., Tancik, M. & Kanazawa, A. pixelNeRF: neural radiance fields from one or few images. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 4578–4587 (2021).
175. Zhang, K., Riegler, G., Snavely, N. & Koltun, V. NeRF++: analyzing and improving neural radiance fields. Preprint at https://doi.org/10.48550/arXiv.2010.07492 (2020).
176. Bylow, E., Sturm, J., Kerl, C., Kahl, F. & Cremers, D. Real-time camera tracking and 3D reconstruction using signed distance functions. *Rob. Sci. Syst.* **2**, 2 (2013).
177. Park, J. J., Florence, P., Straub, J., Newcombe, R. & Lovegrove, S. DeepSDF: learning continuous signed distance functions for shape representation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 165–174 (2019).
178. Kolouri, S., Abbasi, A., Koohpayegani, S. A., Nooralinejad, P. & Pirsiavash, H. Multi-agent lifelong implicit neural learning. *IEEE Signal Process. Lett.* **30**, 1812–1816 (2023).
179. Bortnik, J. & Camporeale, E. Ten ways to apply machine learning in the earth and space sciences. In *AGU Fall Meeting Abstracts* IN12A-06 (2021).
180. Zhang, Y., Bai, Y., Wang, M. & Hu, J. Cooperative adaptive cruise control with robustness against communication delay: an approach in the space domain. *IEEE Trans. Intell. Transport. Syst.* **22**, 5496–5507 (2020).
181. Gao, Y. & Chien, S. Review on space robotics: toward top-level science through space exploration. *Sci. Rob.* **2**, eaan5074 (2017).
182. Bornstein, B. J. et al. Autonomous exploration for gathering increased science. *NASA Tech Briefs* **34**(9), 10 (2010).
183. Swan, R. M. et al. AI4MARS: a dataset for terrain-aware autonomous driving on Mars. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 1982–1991 (2021).
184. Bayer, T. Planning for the un-plannable: redundancy, fault protection, contingency planning and anomaly response for the mars reconnaissance oribiter mission. In *AIAA SPACE 2007 Conference and Exposition* 6109 (2007).
185. Rieke, N. et al. The future of digital health with federated learning. *NPJ Dig. Med.* **3**, 119 (2020).
186. Sheller, M. J. et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* **10**, 12598 (2020).
187. Moor, M. et al. Foundation models for generalist medical artificial intelligence. *Nature* **616**, 259–265 (2023).
188. Bécue, A., Praça, I. & Gama, J. Artificial intelligence, cyber-threats and industry 4.0: challenges and opportunities. *Artif. Intell. Rev.* **54**, 3849–3886 (2021).
189. Buczak, A. L. & Guven, E. A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Commun. Surv. Tutorials* **18**, 1153–1176 (2015).
190. Shaukat, K., Luo, S., Varadharajan, V., Hameed, I. A. & Xu, M. A survey on machine learning techniques for cyber security in the last decade. *IEEE Access* **8**, 222310–222354 (2020).
191. Berman, D. S., Buczak, A. L., Chavis, J. S. & Corbett, C. L. A survey of deep learning methods for cyber security. *Information* **10**, 122 (2019).
192. Kozik, R., Choras, M. & Keller, J. Balanced efficient lifelong learning (B-ELLA) for cyber attack detection. *J. Univers. Comput. Sci.* **25**, 2–15 (2019).
193. Bernstein, D. S., Givan, R., Immerman, N. & Zilberstein, S. The complexity of decentralized control of Markov decision processes. *Math. Oper. Res.* **27**, 819–840 (2002).
194. Goldman, C. V. & Zilberstein, S. Decentralized control of cooperative systems: categorization and complexity analysis. *J. Artif. Intell. Res.* **22**, 143–174 (2004).
195. Melo, F. S., Spaan, M. T. J. & Witwicki, S. J. In *Multi-Agent Systems* (eds Cossentino, M. et al.) 189–204 (Springer, 2012).
196. Vaswani, A. et al. Attention is all you need. In *Advances in Neural Information Processing Systems* Vol. 30 (2017).
197. Khan, S. et al. Transformers in vision: a survey. *ACM Comput. Surv.* **54**, 200 (2022).
198. Bommasani, R. et al. On the opportunities and risks of foundation models. Preprint at https://doi.org/10.48550/arXiv.2108.07258 (2021).
199. Yang, S. et al. Foundation models for decision making: problems, methods, and opportunities. Preprint at https://doi.org/10.48550/arXiv.2303.04129 (2023).
200. Knight, W. OpenAI's CEO says the age of giant AI models is already over. *Wired* https://www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over/ (17 April 2023).
201. Rahwan, I. et al. Machine behaviour. *Nature* **568**, 477–486 (2019).
202. Cath, C. Governing artificial intelligence: ethical, legal and technical opportunities and challenges. *Phil. Trans. R. Soc. A* **376**, 20180080 (2018).

203. Cao, Y. & Yang, J. Towards making systems forget with machine unlearning. In *2015 IEEE Symposium on Security and Privacy* 463–480 (IEEE, 2015).
204. Nick, B. *Superintelligence: Paths, Dangers, Strategies*. (Oxford Univ. Press, 2014).
205. Marr, B. The 15 biggest risks of artificial intelligence. *Forbes* https://www.forbes.com/sites/bernardmarr/2023/06/02/the-15-biggest-risks-of-artificial-intelligence/?sh=309f29002706 (2 June 2023).
206. Bengio, Y. et al. Managing AI risks in an era of rapid progress. Preprint at https://doi.org/10.48550/arXiv.2310.17688 (2023).
207. Wu, C.-J. et al. Sustainable AI: environmental implications, challenges and opportunities. *Proc. Mach. Learn. Syst.* **4**, 795–813 (2022).

Acknowledgements
This material is based on work supported by DARPA under contracts HR00112190132, HR00112190133, HR00112190134, HR00112190135, HR00112190130 and HR00112190136. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA. The authors would like to thank B. Bertoldson, A. Carta, B. Clipp, N. Jennings, K. Stanley, C. Ekanadham, N. Ketz, M. Paravia, M. Petrescu, T. Senator and J. Steil for constructive discussions and comments on early versions of the manuscript.

Author contributions
All authors contributed with insights during brainstorming, ideas and writing the paper. A.S. conceived the main idea and led the integration of all contributions.

Competing interests
P.S. serves as the executive director of Sony AI America and receives financial compensation for this work. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research. All other authors declare no competing interests.

Additional information
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s42256-024-00800-2.

Correspondence should be addressed to Andrea Soltoggio.

Peer review information *Nature Machine Intelligence* thanks Senen Barro, Vincenzo Lomonaco, Xiaoying Tang, Gido van de Ven and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

[1]Computer Science Department, Loughborough University, Loughborough, UK. [2]Computer Science Department, Rice University, Houston, TX, USA. [3]University of Pennsylvania, Philadelphia, PA, USA. [4]ECS Federal, Arlington, VA, USA. [5]Thomas Lord Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA, USA. [6]Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA. [7]Department of Diagnostic and Interventional Imaging, The University of Texas McGovern Medical School at Houston, Houston, TX, USA. [8]The Department of Radiology and Oncology, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. [9]Smart Information Flow Technologies, Minneapolis, MN, USA. [10]Aurora Flight Sciences, Cambridge, MA, USA. [11]Department of Computer Science, Vanderbilt University, Nashville, TN, USA. [12]Massachusetts Institute of Technology, Cambridge, MA, USA. [13]Department of Electrical Engineering, Yale University, New Haven, CT, USA. [14]Department of Computer Science, University of California, Davis, Davis, CA, USA. [15]Department of Diagnostic Radiology and Nuclear Medicine, University of Maryland School of Medicine, Baltimore, MD, USA. [16]Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. [17]Department of Software and Information System Engineering, Ben-Gurion University, Beer Sheva, Israel. [18]University of Massachusetts, Amherst, Amherst, MA, USA. [19]Department of Computer Science, The University of Texas at Austin, Austin, TX, USA. [20]Sony AI America, Sony AI, Austin, TX, USA. [21]Simons Institute, University of California, Berkeley, Berkeley, CA, USA. [22]Department of Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, CA, USA. ✉e-mail: a.soltoggio@lboro.ac.uk