

Cumulative learning

Thórisson, Kristinn R.; Bieger, Jordi; Li, Xiang; Wang, Pei

DOI

[10.1007/978-3-030-27005-6_20](https://doi.org/10.1007/978-3-030-27005-6_20)

Publication date

2019

Document Version

Accepted author manuscript

Published in

Artificial General Intelligence - 12th International Conference, AGI 2019, Proceedings

Citation (APA)

Thórisson, K. R., Bieger, J., Li, X., & Wang, P. (2019). Cumulative learning. In P. Hammer, P. Agrawal, B. Goertzel, & M. Iklé (Eds.), *Artificial General Intelligence - 12th International Conference, AGI 2019, Proceedings* (pp. 198-208). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 11654 LNAI). Springer.
https://doi.org/10.1007/978-3-030-27005-6_20

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Cumulative Learning^{*}

Kristinn R. Thórisson^{1,2}, Jordi Bieger^{1,3}, Xiang Li⁴, and Pei Wang⁴

¹ Center for Analysis and Design of Intelligent Agents, Reykjavik University, Iceland

² Icelandic Institute for Intelligent Machines, Iceland

³ Faculty of Technology, Policy & Management, Delft University of Technology, Netherlands

⁴ Temple University, USA

thorisson@ru.is, j.e.bieger@tudelft.nl,
xiangliAGI@temple.edu, pei.wang@temple.edu

Abstract. An important feature of human learning is the ability to continuously accept new information and unify it with existing knowledge, a process that proceeds largely automatically and without catastrophic side-effects. A generally intelligent machine (AGI) should be able to learn a wide range of tasks in a variety of environments. Knowledge acquisition in partially-known and dynamic task-environments cannot happen all-at-once, and AGI-aspiring systems must thus be capable of *cumulative learning*: efficiently making use of existing knowledge while learning new things, increasing the scope of ability and knowledge incrementally—without catastrophic forgetting or damaging existing skills. Many aspects of such learning have been addressed in artificial intelligence (AI) research, but relatively few examples of cumulative learning have been demonstrated to date and no generally accepted explicit definition exists of this category of learning. Here we provide a general definition of cumulative learning and describe how it relates to other concepts frequently used in the AI literature.

Keywords: Cumulative Learning, Autonomous Knowledge Acquisition, Knowledge Representation, Artificial General Intelligence

1 Introduction

To be autonomous, any learner in the physical world must be able to learn incrementally over time, as it is impossible to be in multiple places at once; equally importantly, one cannot know up front everything that may be relevant in the future (in which case learning would be mostly unnecessary). A learning mechanism that avoids putting acquired experience in silos, through generalization and old-new unification, will always be more parsimonious, and thus more effective, than the alternatives. We consider such *cumulative learning* a hallmark of human cognition, and of central importance to artificial general intelligence (AGI).

The concept of cumulative learning offers an integrated view of numerous cognitive processes that largely have been treated in isolation in the AI literature to date, including, in one form or another, pattern matching, reasoning, continuous information

^{*} Partial funding for this work was provided by Isavia Iceland, IIIM, Reykjavik University, and Delft University of Technology.

acquisition, and old-new integration or unification (which has mostly been ignored outside of [15]). While a fragmented approach to learning may be sufficient for narrow AI research, we consider the current state of fragmentation to be detrimental to AGI research, and call for a more integrated perspective to help the field avoid obscuring important phenomena and slowing down progress towards artificial *general* intelligence. This paper attempts to bring all relevant concerns into one place and set the stage for a long-term vision for cumulative learning in AGI systems.

At the heart of cumulative learning is a process of *unification*: New information enters by default into a process of being integrated with already-acquired knowledge—whether it is in agreement with it or not. This is compression under requirements of *incrementality*, *realtime*,¹ and *generalization*: Replacing incorrect knowledge and extending current knowledge frequently, while generalizing when possible, prepares knowledge to be efficiently applicable to the largest possible class of situations, tasks, topics, and domains—as soon as possible during the learner’s lifetime.

Several aspects of cumulative learning as formulated here² have been covered in the machine learning literature, *but its many necessary-but-not-sufficient features* have invariably been addressed *in isolation*. As any student of systems engineering knows, it is infeasible to join disparate mechanisms, based on incompatible theoretical foundations, into a single coherent system. Due to the lack of a coherent, comprehensive theory of learning, research on this topic in various fields has yielded a number of ontologically inconsistent terms for the various aspects of the phenomenon, and the almost complete ignorance of the importance of incremental knowledge unification. *Always-on* learning has for instance variously appeared under the headings ‘lifelong’, ‘perpetual’, ‘never-ending’, ‘incremental’, ‘online’, and ‘continual’ learning [8,14,19,29,30], most of which only have partial overlap. Other examples of concepts prevalent in the literature of varying relevance to cumulative learning include ‘learning to learn,’ ‘multi-task learning,’ ‘metalearning,’ ‘transfer learning,’ ‘domain adaptation,’ ‘inductive/knowledge transfer,’ ‘knowledge consolidation,’ ‘knowledge-based inductive bias,’ ‘context-sensitive learning,’ ‘catastrophic forgetting/interference,’ and ‘semi-supervised learning’ [5,10,17,23]. Few systems have proposed to address the full scope of cumulative learning as formulated here. Two systems that explicitly have, and presented empirical evidence of progress towards it, are the Non-Axiomatic Reasoning System (NARS) [25,26] and Auto-Cataytic Endogenous Reflective Architecture (AERA) [15,16].

In addition insufficient focus on old-new unification, few of the above concepts have been conceived in the context of (artificial) general intelligence and are thus in one or more aspects at odds with the larger, more complete picture of learning that we find needed for AGI. Here we attempt to present a coherent picture by ‘defragmenting’ the conceptual space surrounding learning, painting instead a coherent picture more suited as a step towards a theory of cumulative learning.

¹ Unification must happen frequently, relative to the learner’s lifetime, lest we’d be hard-pressed to call it ‘cumulative.’

² The term itself has appeared in the AI literature with some overlap of its sense here (cf. [6,7,2]), as well as in AGI research (cf. [22]).

2 Dimensions of (Cumulative) Learning

Learning is necessary for goal achievement in a changing, novel environment. All learning machines, whether natural or artificial, are limited by the time and energy they have available; the outermost constraint on any learning mechanism is the assumption of insufficient knowledge and resources (AIKR) [27]. However, there is a large number of ways to interpret these constraints when implementing learning mechanisms, and thus there are numerous dimensions along which any learning ability may vary. We have identified 14 dimensions whose settings determine the performance characteristics of any particular cumulative learning implementation. These naturally form three sets: (1) *Memory management*, (2) *temporal capacity and granularity*, and (3) *generality*. In each group there are between four and six different dimensions that we will now outline. While these are not perfectly orthogonal to each other (which would require a proper theory of learning), the breakdown allows us to better place prior work in the context of the present focus. Note that our focus in this paper is not so much on learning *methods* but primarily on externally measurable factors and characteristics of the cumulative learning process as a whole, and related learner performance characteristics.

- [A] **Memory Management.** Operational characteristics of processes related to memory and knowledge management. These can of course also be learned, i.e. improved with experience. Having to do with quality, these range from *catastrophic* at one end to *highly effective* at the other.
- (a) **Storage:** Storing relevant aspects of experience in memory.
 - (b) **Remembrance:** Bringing relevant knowledge to bear on a task or problem.
 - (c) **Forgetting:** Removing the least relevant and necessary knowledge, if needed.
 - (d) **Compression:** “Cleaning up” knowledge in ways that can improve the learner in some way, w.r.t. storage, forgetting, remembrance, generality, etc.
 - (e) **Old-New Unification:** Integrating new information with existing knowledge so that it becomes more coherent and parsimonious.
 - (f) **Defeasibility:** Replacing less correct, less useful and/or less detailed knowledge with new more correct, more useful and/or more detailed knowledge.
- [B] **Temporal Capacity & Granularity.** When and how the learner can accept new information. This group contains four important characteristics that define temporal measures of cumulative learning:
- (g) **Concurrent capacity:** How many things³ can be learned concurrently.
 - (h) **Consecutive capacity:** How many things can be learned consecutively.
 - (i) **Temporal granularity** of information **acceptance/storage**.
 - (j) **Temporal granularity** of old-new information **unification**.⁴

The range of dimensions **Bg** and **Bh** starts with *one thing once* at the lower end, meaning that a single learner can only learn one thing at one (particular) time, and

³ A “thing” can be a task, environment, goal, domain, phenomenon, process, etc.—it does not matter so much here which, as long as there is some way to compare systems on these features.

⁴ ‘Learning’ here means the acquisition of knowledge applicable to achieving goals the learner might face, now or in the future; this view does not address “non-actionable knowledge”.

extends towards *many things at any time, concurrently/simultaneously and/or consecutively/sequentially* at the other end, meaning the learner can at any time learn new things, no matter how large or small (ignoring learning time and materials). **Bi** and **Bj** at the lower end range from *a single two-step learn-then-apply pair* (e.g. artificial neural nets), to *concurrently and/or consecutively and continuously* (non-discretized) at the other end.

- [C] **Generality** of the learning, with respect to task, goal, domain, etc. These parameters range from *one* at the lower end, to *any* at the other.
- (k) **Data Flexibility**: Flexibility in the kind of data that learner can accept (as dictated by cognitive – not sensing – capabilities).
 - (l) **Knowledge Flexibility**: Flexibility in what knowledge is leveraged.
 - (m) **Knowledge Transfer**: Using knowledge acquired for one purpose in one context to other purposes and contexts.
 - (n) **Learning to Learn**: Using acquired knowledge to improve learning ability.
 - (o) **Inverse Defeasibility**: New information improves existing knowledge. The more generally a learner can do this (i.e. the less directly the new information is related to current knowledge) the better a learner it is.

3 Functions of Cumulative Learning

A cumulative learner in our conceptualization is a learning controller [22] that, guided by one or more top-level internalized goals (or drives), implements a cumulative modeling process whereby regularities are recursively extracted from the learner’s experience (of self and environment) to construct integrated models useful for achieving goals [3,22]. The collection of models form a unified body of knowledge that can be used, by a set of additional and appropriate management processes (see **Aa-f** above), as the basis for making predictions about, and achieving goals with respect to, an environment, and that can be used to improve future learning—in speed, quality, efficiency, or all of these [28]. At the risk of oversimplification, a compact definition of cumulative learning might read something like “using a unified body of knowledge to continually and recursively integrate new information from many contexts into that body.” A learner whose models capture in this way an actionable description of measurable, verifiable entities in the environment and their relations, and tends over time towards the pragmatically simplest (interrelated) set of such descriptions, is in some sense an ideal cumulative learner.

We will now turn to the central features of this conceptualization of cumulative learning in light of notable related work. We emphasize that our interest in cumulative learning is limited to a learner that *embodies and unifies all* of the following learning-related functions *in a single coordinated system, throughout its lifetime*, as is necessary for making progress towards AGI.

3.1 Temporal Capacity & Granularity

An important dimension of learning concerns how and when the learner is open to accepting, storing and integrating new information into its knowledge base.

Learning Multiple Things. A cumulative learner must, by definition, be able to learn multiple things – tasks, goals, environmental factors, techniques, rules of thumb, generalizations, modes of reasoning, etc. – cumulatively over time: It must not be restricted to a single function, task, domain or phenomenon. Aspects of this capability have been studied under the term ‘multitask learning’ (MTL) [4] where the learner learns multiple tasks concurrently (cf. **Bg**). MTL assumes the input representation for each task is the same, and concurrent learning requires predefined and pre-programmed knowledge of the number of tasks, and (ideally) access to a data set where each input is associated with a target output label (in the supervised learning setting for which it was conceived). Fei et al. [7] use the term ‘cumulative learning’ to describe a variation of this type of MTL, where tasks are added one after the other (cf. **Bh**).

MTL can be extended to control tasks in a reinforcement learning setting [21] by assuming the tasks are encountered consecutively (rather than assuming a single agent simultaneously acting in multiple task-environments; cf. **Bh**). In this setting MTL research often makes use of hierarchical reinforcement learning (HRL), which also involves learning multiple (sub)tasks that together constitute a top-level task. When the top-level task is removed, leaving just the subtasks, this is closely related to both multitask learning and multi-objective reinforcement learning (MORL) where an agent has multiple active goals. This kind of process can happen organically in NARS [25].

An ideal cumulative learner should be capable of learning multiple things both concurrently and consecutively, as appropriate, without constraints on the order of encountered phenomena and task revisitation. NARS [24] can accept input data of any content at any time (cf. **Bi, Bj**), as long as they can be expressed in a format that the system can recognize (cf. **Ck, Cl**). This means NARS has the ability to solve any problems expressible in its language instead of being limited to a specific type of problems. When solving a problem, all accumulated evidence matters, though different pieces may contribute differently (cf. **Ab, Ae, Af**). In AERA [15] any digital data can become a source of learning (cf. **Ck, Cl**), and the simplest basis on which it can learn is correlation between data types, variables or values (cf. **Aa**). To learn, AERA must from the outset have a drive in its seed that references one or more observable environment variables.

Always-On Learning. An ideal cumulative learner can learn at any time—there are no pre-designated periods where its learning is turned off, no required dedicated special training phase (although it is of course not prevented), and the learning process does not converge to an attractor (cf. **Bi, Bj**). Thus, learning occurs perpetually throughout the operational lifetime of a cumulative learner.

Lifelong machine learning (LML) [19,6], continual learning [18], perpetual learning [30] and never-ending learning [14] all focus on sequentially learning an unknown number of tasks.⁵ As a result, learning in these settings never truly ends. However, this does not necessarily mean that learning is always on. For instance, Zhang’s perpetual learner [30] only enters a new learning phase when a “learning stimulus” is encountered (i.e. an inconsistency, anomaly or surprise) during each (learning-free) application

⁵ While the terms ‘lifelong learning’ and ‘lifelong machine learning’ are not always used entirely consistently, they can be considered approximately interchangeable with ‘perpetual learning’ and ‘never-ending learning,’ respectively.

phase. Furthermore, many lifelong learners consider learning on the current task “done” before starting the new one and it is typically not clear when the learned knowledge is supposed to be applied (and what can be learned from that application), suggesting that even here there is a separation between training/learning and application phases.

The temporal granularity at which incoming information can be accepted, stored and added to the knowledge base are important dimensions of learning (cf. **Bi**, **Bj**). While many ML systems can only learn in a single designated phase at the beginning of their lifetime followed by a phase in which this knowledge is applied, other systems can alternate between these modes (e.g. Zhang’s perpetual learner [30]), while yet others learn constantly with or without explicit learning/application phases (e.g. NARS [24] and AERA [15]). The rate at which new information can be accepted and stored, and the rate at which it can be usefully unified into the knowledge base, are separate dimensions.

Assessing temporal granularity of a learner involves examining how much information it needs before learning can occur. *Offline* or *batch learning* assumes constant on-demand access to all data, no restrictions on time and space for training, and a fixed (often *i.i.d.*) distribution from which the data is pulled, while *online* or *incremental learning* removes these assumptions [8]. In online/incremental learning information is encountered sequentially and there are often restrictions placed on the ability to revisit old data. In the most extreme case, upon encountering some new datum d the learner’s model m must be updated based only on m and d , without considering any previously encountered data. A continuum of incrementality could be considered based on how much previous data can be used to update m , where offline/batch learning is at the other extreme because it uses all data.

Incrementality in LML can be evaluated at multiple levels: While e.g. tasks are often encountered sequentially, and data from previous tasks may or may not be available, it is often the case that each individual task is trained offline when it is encountered. Online learning is common in forecasting, sequence prediction, and sequential decision making. Many reinforcement learning algorithms learn online (e.g. Q-learning), although other algorithms (e.g. policy gradient) and function approximations (e.g. using deep learning) may require batches of data.

3.2 Memory Management

For an implemented system, neither memory nor computation speed is infinite [27]. This means all learners must make choices on what knowledge can and should be retained (cf. **Aa**). Systems that cannot forget will inevitably run into memory limits at some point and break down, or demand human intervention, either of which are sub-optimal because processing an ever increasing amount of knowledge will become prohibitive due to the limitation on computation speed (cf. **Ab**, **Ac**).

When learning a new task causes forgetting of critical or all parts of previously learned tasks, this is called *catastrophic forgetting* [11]. Workarounds include e.g. “freezing” of knowledge obtained for previously encountered tasks, and retaining training data to engage in task rehearsal (i.e. continuously retraining on old tasks as well as new ones), but this runs into aforementioned limits of space and time. An important challenge to address in cumulative learning is thus the stability-plasticity balance [13],

wherein sufficient plasticity is needed to incorporate new knowledge while retaining old knowledge requires sufficient stability.

Forgetting sensibly is bound to involve several processes, such as replacing wrong or worse knowledge with correct or better knowledge, respectively, whenever possible (cf. **Af**). There should be multiple ways of compressing the knowledge (with or without loss; cf. **Ad**)—induction (generalization) is one way to do so, forgetting permanently is another one (based on empirically-evaluated usefulness). Numerous combinations of various mechanisms are possible, achieving various trade-offs between memory requirements, applicability, manageability, and so on. In addition to selective forgetting, AERA’s rewriting rules reduces redundancies and storage requirements through increased generality whereby values are replaced with variables coupled with ranges [15]. In NARS, forgetting has two related senses: (1) relative forgetting: decrease priority to save time, (2) absolute forgetting: remove from memory to save space and time [27].

3.3 Generality

The last set of learning dimensions considered here concerns the generality and generalization ability of the learning system. Ideal cumulative learners can accumulate knowledge of any type and in any situation, and generalize it for use in both unseen future contexts and previously encountered situations. As with the other dimensions, the focus here is not on learning methods, i.e. how generality is achieved, but rather on externally measurable characteristics of cumulative learners and performance.

Domain-, Task- & Goal-Generality. A domain-general (domain-independent) cumulative learner will model any relevant experience, including its own sensors, the quality of data they produce (in relation to other sensors), as well as the quality of data acquired from outside sources (cf. **Ck, Cl**), and even its own cognitive processes. An ideal artificial cumulative learner, in our conceptualization, can therefore acquire knowledge and skills through both experience [20] and explicit teaching [3]. Goal-generality means that knowledge and goal(s) are not fused together (in particular situations and constraints) but can be re-purposed when task- and domain-related parameters change [9].

It is worth pointing out that paradigms like transfer learning, MTL and LML tend to focus on the task as a distinct unit (cf. **Bi, Bj**): It is assumed that tasks are explicitly separated from the point of view of the learner, who is typically notified when learning on a new task starts, or of the task that should currently be performed. In the general case of the real physical world, task boundaries are not this clear. (Is playing tennis against well-known tennis player Roger Federer a different task than playing against Rafael Nadal? What about playing doubles? Or against a child? What about playing squash or badminton?) Correctly recognizing contexts and knowing what prior knowledge to bear (and how) is a key part of the challenge that cumulative learning solves: Boundaries between tasks and domains for autonomous learners in the real world are inexplicit. Animals learn continuously, cumulatively adding new knowledge to their current knowledge, as needed. NARS [24] accepts input data and task of any contents, as far as they can be expressed in a recognizable format (cf. **Ck, Cl**). AERA [15] is data-general as its learning methods are data-agnostic (while its learning is not) (cf. **Ck, Cl**).

Unlike transfer learning, with its explicit focus on the learning period itself, cumulative learning assumes a continually running process of unification—irrespective of how or whether the new knowledge can be, or was, useful in learning the new information (cf. **Bi, Bj**). An extreme case of this is using analogies to deepen or broaden knowledge of a set of phenomena. In NARS, for instance, learning involves not only tasks but also effective (re-)organization of knowledge, without respect to specific problems, so that it may later be used on *any* relevant task [28] (cf. **Ae, Af**). The idea of such meta-learning (‘super-task learning’ or ‘task-free learning’) is naturally only a challenge in a context where multiple things are actually learned, and has only recently received some attention [1,27]. In AERA [15] models are by themselves general in that they are not attached to any particular task (this is always computed on a case-by-case basis on the fly), and each model is thus in principle applicable to any part of any task, as long as its preconditions are met (cf. **Ck, Cl**).⁶

Knowledge Transfer. Cumulative modeling, to achieve effective compression and old-new unification for any new context or situation, needs to find ways of dealing with similar input at different points in time, and note its similarity and differences, so that old knowledge acquired at time t_1 for situation S_1 can be successfully used for a new situation S_2 at time t_2 . This can be done by e.g. making analogies [24] (cf. **Cm**). New information should be integrated with old information, at a low level of detail (as low as possible, in each case), producing a growing set of interrelated (fine-grained) models of the task-environment [15] (cf. **Ae**).

Similarly, the goal of transfer learning and domain adaptation is to use knowledge obtained in a set (typically of size one) of previously learned source tasks in order to facilitate learning and/or performance on a target task [17] (cf. **Cm, Cn**). Perhaps more generally, it deals with the situation where (some or all of) the training is obtained in a situation different from the one in which it is to be applied. Making use of existing knowledge (‘inductive bias’) can enable faster learning from one or a few observations that would otherwise not contain enough information (‘one-shot’ or ‘few-shot’ learning), possibly even without ever direct observation (‘zero-shot learning’) [12].

To make knowledge transfer between tasks and situations positive (helping instead of hurting learning and performance), it is important to consider what, when, and how relevant knowledge is transferred. Most work to date has focused on “how,” while relevance of prior knowledge is already assumed, and assumed that most transfer happens right before the learner starts learning a target task. Work on task similarity and transferability is rarer, as is the question of when to transfer. An ideal cumulative learner will always treat new information in a way that makes it generally applicable to future tasks, so there is no explicit knowledge transfer step or stage—just the future application of the most relevant available knowledge in each instance. This is how knowledge transfer and learning works in NARS [24] and AERA [15]. Furthermore, at the present time what we might call “forward transfer” – the effect of current learning on future learning – is considered more important than “backward transfer” (the effects that learning the

⁶ Still, if some models are often used together they may be compiled for faster future use, which may be comparable to detecting “tasks” that are meaningful to the learning system).

new task has on the ability to perform the previously learned tasks). In practice, backward transfer in much of current machine learning is typically extremely negative, as catastrophic interference/forgetting [11] occurs where the previous tasks are forgotten almost entirely or performance drops dramatically (cf. **Bc**).

4 Conclusions

Artificially generally intelligent (AGI) systems will need to handle unknown dynamic environments, where required information cannot be known fully up front and many skills must be acquired. Cumulative learning, as we conceive of it, is important for AGI for numerous reasons, including: (1) Knowledge is created incrementally, matching the needs of partially-known, changing environments, (2) knowledge is built up and improved in small increments, avoiding pitfalls of catastrophic forgetting and errors, (3) new knowledge is immediately available to the learner, and (4) knowledge consisting of fine-grained (low-level) explicit models provides explicitness necessary for comparing, managing, reasoning, etc. To be useful for AGI systems these skills must *all exist in a unified manner in one and the same learner*. In this paper we have tried to clarify why and how the various aspects of cumulative learning relate to key AGI requirements, and place it in the context of prior work. More work is clearly needed to realize true artificial cumulative learning in a single system on par with that found in nature. The systems developed by the authors, NARS [25] and AERA [15], demonstrate some important steps in this direction by bringing several of its features together in single unified systems.

References

1. Aljundi, R., Kelchtermans, K., Tuytelaars, T.: Task-free continual learning. CoRR (2018)
2. Baldassare, G., Mirolli, M., Mannella, F., Caligiore, D., Visalberghi, E., Natale, F., Truppa, V., Sabbatini, G., Guglielmelli, E., Keller, F., others: The IM-CLeVeR project: Intrinsically motivated cumulative learning versatile robots. In: 9th International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems. pp. 189–190 (2009)
3. Bieger, J.E., Thórisson, K.R.: Task analysis for teaching cumulative learners. In: Proc. Artificial General Intelligence. pp. 21–31. Springer International Publishing (2018)
4. Caruana, R.A.: Multitask connectionist learning. In: Proceedings of the 1993 Connectionist Models Summer School. pp. 372–379 (1993)
5. Chapelle, O., Schölkopf, B., Zien, A.: Semi-supervised learning. Adaptive Computation and Machine Learning, MIT Press, Cambridge, MA (2006)
6. Chen, Z., Liu, B.: Lifelong Machine Learning. Morgan & Claypool Publishers (2016)
7. Fei, G., Wang, S., Liu, B.: Learning cumulatively to become more knowledgeable. In: Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1565–1574. KDD '16 (2016)
8. Fontenla-Romero, Ó., Guijarro-Berdias, B., Martínez-Rego, D., Pérez-Sánchez, B., Petriro-Barral, D.: Online machine learning. Efficiency and Scalability Methods for Computational Intellect pp. 27–54 (2013)
9. Hammer, P., Lofthouse, T.: Goal-directed procedure learning. In: International Conference on Artificial General Intelligence. pp. 77–86. Springer (2018)
10. Jiang, J.G., Su, Z.P., Qi, M.B., Zhang, G.F.: Multi-task coalition parallel formation strategy based on reinforcement learning. Acta Automatica Sinica **34**(3), 349–352 (2008)

11. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**(13), 3521–3526 (2017)
12. Lake, B., Salakhutdinov, R., Gross, J., Tenenbaum, J.: One shot learning of simple visual concepts. In: *Proceedings of the Annual Meeting of the Cognitive Science Society*. vol. 33 (2011)
13. Mermillod, M., Bugaiska, A., Bonin, P.: The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects. *Front. Psychol.* **4** (2013)
14. Mitchell, T., Cohen, W., Hruschka, E., Talukdar, P., Yang, B., Betteridge, J., Carlson, A., Dalvi, B., Gardner, M., Kisiel, B., Krishnamurthy, J., Lao, N., Mazaitis, K., Mohamed, T., Nakashole, N., Platanios, E., Ritter, A., Samadi, M., Settles, B., Wang, R., Wijaya, D., Gupta, A., Chen, X., Saparov, A., Greaves, M., Welling, J.: Never-ending learning. *Commun. ACM* **61**(5), 103–115 (2018)
15. Nivel, E., Thórisson, K.R., Steunebrink, B.R., Dindo, H., Pezzulo, G., Rodriguez, M., Hernandez, C., Ognibene, D., Schmidhuber, J., Sanz, R., Helgason, H.P., Chella, A., Jonsson, G.K.: Bounded Recursive Self-Improvement. Technical RUTR-SCS13006, Reykjavik University Department of Computer Science, Reykjavik, Iceland (2013)
16. Nivel, E., Thórisson, K.R., Dindo, H., Pezzulo, G., Rodriguez, M., Corbato, C., Steunebrink, B., Ognibene, D., Chella, A., Schmidhuber, J., Sanz, R., Helgason, H.P.: Autocatalytic Endogenous Reflective Architecture. Technical RUTR-SCS13002, Reykjavik University School of Computer Science, Reykjavik, Iceland (2013)
17. Pan, S.J., Yang, Q.: A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering* **22**(10), 1345–1359 (Oct 2010). <https://doi.org/10.1109/TKDE.2009.191>
18. Ring, M.B.: CHILd: A first step towards continual learning. *Machine Learning* **28**(1), 77–104 (1997)
19. Silver, D.L., Yang, Q., Li, L.: Lifelong Machine Learning Systems: Beyond Learning Algorithms. In: *AAAI Spring Symposium: Lifelong Machine Learning* (2013)
20. Steunebrink, B.R., Thórisson, K.R., Schmidhuber, J.: Growing recursive self-improvers. In: *Proceedings of Artificial General Intelligence*. pp. 129–139 (2016)
21. Taylor, M.E., Stone, P.: Transfer learning for reinforcement learning domains: A survey. *J. Mach. Learn. Res.* **10**, 1633–1685 (2009)
22. Thórisson, K.R., Talbot, A.: Cumulative learning with causal-relational models. In: *Artificial General Intelligence*. pp. 227–237. Springer International Publishing, Cham (2018)
23. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. *Artificial Intelligence Review* **18**(2), 77–95 (2002)
24. Wang, P.: *Rigid Flexibility: The Logic of Intelligence*. Springer, Dordrecht (2006)
25. Wang, P.: From NARS to a thinking machine. *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms* **157**, 75–93 (2007)
26. Wang, P.: *Non-Axiomatic Logic: A Model of Intelligent Reasoning*. World Scientific Publishing, Singapore (2013)
27. Wang, P.: *Non-Axiomatic Logic: A Model of Intelligent Reasoning*. World Scientific, Singapore (2013)
28. Wang, P., Li, X.: Different conceptions of learning: Function approximation vs. self-organization. In: *Proceedings of Artificial General Intelligence* (2016)
29. Zhan, Y., Taylor, M.E.: Online Transfer Learning in Reinforcement Learning Domains. *arXiv preprint arXiv:1507.00436* (2015)
30. Zhang, D.: From one-off machine learning to perpetual learning: A step perspective. *IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (2018)