IWSSL: Introduction To This Volume

Kristinn R. Thórisson THORISSON@{RU.IS, IIIM.IS}
Center for Analysis and Design of Intelligent Agents, Reykjavik University, Iceland
Icelandic Institute for Intelligent Machines, Reykjavik, Iceland

Editor: Kristinn R. Thórisson

The papers in this collection were presented at the third annual **International Workshop on Self-Supervised Learning** (IWSSL-22), held at Reykjavik University in Reykjavik, Iceland, on July 28^{th} and $29,^{th}$ 2022. Weighing in together at a total of 127 pages, these 10 papers of exceptional quality describe leading AI research from Europe and the USA, both theoretical and applied. The challenges they focus on, and the many foundational ideas proposed to address them, are sure to be at the center of all significant advances in AI research in the coming decades.

1. Our Automated Future

There is little doubt that society will increasingly be shaped by advances in the research and development of artificial intelligence (AI). We are on the cusp – at the *very* beginning – of a century-long transformation to a radically automated future. As AI researchers, our goal should be to make machines smarter, more useful, safer, and more powerful. That's what all the papers in this volume do: They propose, apply, or expand on methods, approaches and systems that in one way or other generate transparent, explainable knowledge from scratch, autonomously, which they then employ in their subsequent autonomous operations.

2. Why Self-Supervised Learning?

In prior introductions to IWSSL PMLR volumes (Robertson et al., 2021) we have explained how the standard methodology in AI has historically been based on humans providing machines directly with knowledge—whether the systems are symbolic, like the expert systems of the 1970s or sub-symbolic like the deep neural networks of the 2020s, an ontology, rules, and key operating principles are defined by hand and spoon-fed as "ground truth." Systems built this way tend to be suited only to a limited set of problems, significantly restrained in their applicability to other domains—even similar ones.

What about machines that learn autonomously—machines that can continue to learn after they leave the lab, re-evaluate their own knowledge and revise it, improving it of their own accord in light of new evidence? How would we make such machines, that in a sense

© 2022 K.R. Thórisson.

program themselves? Addressing these questions possibly calls for more radical approaches than usually employed in AI, calling for a revision of our standard research methodology and philosophy largely inherited from computer science. In fact, as recounted in the paper by Thórisson and Minsky (2022), past presidents of the AAAI Society have repeatedly pointed as much out in their presidential addresses (cf. McCarthy, 1983; Waltz, 1999), albeit to limited impact. AAAI president John McCarthy wrote in 1993 that *"Too few people are doing basic research in AI relative to the number working on applications."* (McCarthy, 1983, p. 5). For whatever reason, that seems to be the current situation, still, today. We see the work in this volume as tipping back that balance, if even just a little bit.

3. The Papers in IWSSL-22

As in the prior three years of the IWSSL, the papers in this volume represent the state of the art in a growing field of research aiming to build systems that can learn without human intervention, with little hard-wired domain knowledge, as would a new-born child or animal. The papers cover fundamental topics that beg for one or more breakthroughs to move the field of AI notably forward. While some of the topics are increasingly identified as important by contemporary AI researchers, attempts at addressing them have been few and far between, and based on approaches dating as far back to the middle of last century, among them artificial neural networks, Bayesian inference, and formal logic, all of which are known to harbor serious limitations for achieving general learning, reasoning and autonomy. At their core, good methodologies must rest on strong theories, and of course, developing good approaches requires taking the long view and dare to address deep challenges. This is one of the hallmarks of the work in this volume.

In the first paper, titled *The Future of AI Research: Ten Defeasible 'Axioms of Intelligence'*, Thórisson and Minsky (2022) present a list 10 disprovable claims about intelligence, intended to motivate young researchers to look beyond the limited scope of contemporary AI to address some of the grand challenges that singularly belong in this field. They point out topics that contemporary methodologies are incapable of addressing, such as existential autonomy, general-purpose control, cumulative learning, and many more, many of which are also the topic of other papers in this collection.

All the papers present in some way one or more new ideas or methodology to constructing intelligent systems. Several of the papers have an implicit or explicit constructivist (cf. Drescher, 1989; Thórisson, 2012) angle, perhaps most obvious in Thórisson and Minsky (2022) (*The Future of AI Research: Ten Defeasible 'Axioms of Intelligence'*), Georgeon et al. (2022) (*Simultaneous Localization and Active Phenomenon Inference (SLAPI)*) and Komrusch and Minsky (2022) (*Symbolic Guidance for Constructivist Learning by Neural Mode*), but also alluded to in Zhou et al. (2022) (*Novel Primitive Decompositions for Real-World Physical Reasoning*), Sheikhlar et al. (2022) (*Explicit Analogy for Autonomous Transversal Learning*) and Hammer (2022) (*Reasoning-Learning Systems Based on Non-Axiomatic Reasoning System Theory*).

The kind of decomposition proposed by Zhou et al. (2022), in their paper *Novel Primitive Decompositions for Real-World Physical Reasoning*, is concerned with cognitive representations for enhancing an agent's ability for self-supervised language learning. Seeing language as just another skill that can be learned, their approach involves creating a representational substrate that is language-independent, including image schemas and "mental motion pictures," that can address object permanence, movement, and spatial relationships.

In his paper *Intelligence: From Definition to Design* Wang (2022) provides a non-technical description of the ideas behind the author's non-axiomatic reasoning model of intelligence, and explains some of the major design decisions of the model, comparing it to well-known approaches in AI.

Seeing the usefulness of analogies as a way to learn about novel phenomena, in a paper titled *Explicit Analogy for Autonomous Transversal Learning* Sheikhlar et al. (2022) present a theory of analogy and show how it can be implemented in the runtime learning system AERA by addressing three necessary prerequisites for autonomous analogy creation.

In the paper *Symbolic Guidance for Constructivist Learning by Neural Model* Komrusch and Minsky (2022) describe how an artificial neural network with internal feedback can be used to propose proper actions to be taken, and predict the results of those actions, based on an autonomously-learned internal world model that supports reasoning.

A paper titled *Simultaneous Localization and Active Phenomenon Inference (SLAPI)* by Zhou et al. (2022) addresses the localization tasks that a mobile robot may face, and how it can actively infer the existence and properties of novel phenomena present in its surrounding environment, what they call the "SLAPI problem." Their robot contains a number of sensors that feed the creation of knowledge about a novel phenomenon's affordances.

Hammer (2022), in his paper titled *Reasoning-Learning Systems Based on Non-Axiomatic Reasoning System Theory*, describes learning can be facilitated through autonomous reasoning, building knowledge of its own experience incrementally over time. The author compares the learning process on various tasks to reinforcement learning methods.

The three last papers in the collection are focused on applications and are authored from the viewpoint of industrial needs. In the first of these, titled *Neurosymbolic Learning on Activity Summarization of Video Data* at Leela AI, Komrusch et al. (2022) describe how the output of an object- and pose-detecting artificial neural network can be fed as input to a symbolic learning system that then automatically corrects the object and pose data and incrementally trains and improves the neural network. They also show how symbolic learning techniques can improve action detection when given example ground truths by humans. The second of these papers, *The Holon System: Artificial General Intelligence as 'Work on Command,'* Steunebrink et al. (2022) at General Systems Theory, provides a pragmatic definition of general intelligence and presents their highly novel system, which is designed and implemented based on this definition. They explain why deep neural networks cannot meet this criterion for *general* intelligence. The last paper in the volume is by Latapie et al. (2022) at Cisco Systems' Inc. Emerging Technologies & Incubation division, titled *Hybrid AI for IoT Actionable Insights & Real-Time Data-Driven Networks*, describing several experimental projects, built in collaboration with their academic partners (Pei Wang and Patrick Hammer, to name just two), where hybrid neuro-symbolic systems are deployed in industrial scenarios.

References

Gary L. Drescher. *Made-up minds: a constructivist approach to artificial intelligence.* PhD thesis, Massachusetts Institute of Technology, 1989.

Olivier Georgeon, Juan R. Vidal, Titouan Knockaert, and Paul Robertson. Simultaneous localization and active phenomenon inference (SLAPI). *Proceedings of Machine Learning Research*, 192:77–88, 2022.

Patrick Hammer. Reasoning-learning systems based on Non-Axiomatic Reasoning System theory. *Proceedings of Machine Learning Research*, 192:89–107, 2022.

Steve Komrusch and Henry Minsky. Symbolic guidance for constructivist learning by neural model. *Proceedings of Machine Learning Research*, 192:63–76, 2022.

Steve Komrusch, Sanket Bhave, Mridul Banik, and Henry Minsky. Neurosymbolic learning on activity summarization of video data. *Proceedings of Machine Learning Research*, 192: 108–119, 2022.

Hugo Latapie, Mina Gabriel, and Ramana Kompella. Hybrid AI for IoT actionable insights & real-time data-driven networks. *Proceedings of Machine Learning Research*, 192:127–131, 2022.

John McCarthy. Artificial intelligence needs more emphasis on basic research: President's quarterly message. In *AI Magazine*, volume 4, page 5, 1983.

Paul Robertson, Kristinn R. Thórisson, Milan Minsky, and Henry Minsky. IWSSL: Introduction to this volume. *Proceedings of Machine Learning Research*, 159:1–4, 2021.

Arash Sheikhlar, Kristinn R. Thórisson, and Jeff Thompson. Explicit general analogy for autonomous transversal learning. *Proceedings of Machine Learning Research*, 192:48–62, 2022.

Bas Steunebrink, Jerry Swan, and Eric Nivel. The Holon system: Artificial general intelligence as 'work on command'. *Proceedings of Machine Learning Research*, 192:120–126, 2022.

Kristinn R. Thórisson. A new constructivist AI: from manual methods to self-constructive systems. In *Theoretical Foundations of Artificial General Intelligence*, pages 145–171. Springer, 2012.

Kristinn R. Thórisson and Henry Minsky. The future of AI research: Ten defeasible 'axioms of intelligence'. *Proceedings of Machine Learning Research*, 192:5–21, 2022.

David Waltz. The importance of importance: AAAI-98 presidential address. In *AI Magazine*, volume 20, pages 18–35, 1999.

Pei Wang. Intelligence: From definition to design. *Proceedings of Machine Learning Research*, 192:35–47, 2022.

Mackie Zhou, Bridget Duah, and Jamie C. Macbeth. Novel primitive decompositions for real-world physical reasoning. *Proceedings of Machine Learning Research*, 192:22–34, 2022.