KI - Künstliche Intelligenz (2019) 33:205–207
https://doi.org/10.1007/s13218-019-00604-2

EDITORIAL

Bridging the Prototype Gap: On the Evolution of Ugly Ducklings

Anni-Yasmin Turhan[1]

Dear readers,

At a recent workshop on formal logic and reasoning one of the most remarkable sessions was, in fact, an industry panel. The scientific community attending this workshop mainly investigates formal problems on logical reasoning by methods from theoretical computer science. The algorithms investigated and developed by this community are, if at all, implemented in prototype systems. In the industry panel the representatives from companies employing the logical formalisms and the systems developed by said community reported on the obstacles encountered in practice when using these systems. It turned out that (besides other things) new functionality beyond the core functionality of reasoning is needed. In particular, there was need for services that were so far not catered for or were not even anticipated by the scientific community. Furthermore, the report on industry use-cases was surprising as it revealed that some of the problem instances encountered in practice were much more complex than assumed, demanding new expressive means from the logic in use and thus posing new research questions. Some other instances of the industry use-cases were simply challenging for the implemented systems by their sheer size and not so much by their complex structure. Both kinds of instances are posing new requirements to the implementations of such systems and, more remarkably, also to theoretical research.

This points to a more general problem. The development of scientific prototype systems is often being carried out at universities or other academic institutions. Here, most of the times a prototype is implemented in the course of a Ph.D. project. Once the doctoral student has attained her or his degree after a few years, the system is rarely maintained much longer. Furthermore, the user support for such systems in terms of system descriptions, documentation or even consulting is very limited. This comes at no surprise as providing these is a time consuming and non-trivial task that is not immediately necessary or profitable for attaining a degree or to excel in one's scientific community. Even worse, incremental improvements on such systems are hard to sell to the scientific community and have hardly a "market place". Thus the ugly ducklings hatched by academia rarely evolve into fully-fledged swans. Nevertheless, to develop reliable and usable systems requires a proliferation of versions that essentially implement similar, if not the same core functionality. Early exposure of systems developed by academia to conditions outside of their nursery should not so much be a question of sink or swim, but rather of co-evolution with the habitat "out there". On the one hand the feedback loop from industry or from groups that apply such systems is essential as it can guide the development and refinement of prototype systems. On the other hand, the transfer of knowledge on the competent use of new AI techniques is expedient to their success. While this may also be the case for software systems in general, it is even more important for AI systems as the introduction of entirely new techniques requires strong transfer of knowledge to the users of such systems. Without such a transfer the new techniques may be applied in a naive way, leading unsatisfactory results or performance, which in turn can hamper the adoption of the new techniques. The development of successful AI systems is essentially a boots-strapping process spanning

✉ Anni-Yasmin Turhan
Anni-Yasmin.Turhan@tu-dresden.de

[1] TU Dresden, Dresden, Germany

several years and requiring the full loop to and from users of such systems beyond academia.

Another observation is that several scientific communities have serious difficulties to find meaningful benchmarks for their systems. Often out of a lack of alternatives, benchmarks are generated artificially and do not necessarily reflect the conditions present in practical applications. Consequently, the prototype systems are optimized towards goals that need not be realistic or even useful for industrial applications. Here, benchmarks drawn from industrial use-cases would be very valuable. Such benchmarks are certainly not easy to provide in the presence of industrial secrets and data protection rules in general, but the need for them in order to facilitate the co-evolution of systems developed by academia and industry is clearly there.

It remains a severe obstacle for developers of academic prototypes which are about to evolve into mature systems, that improvements on their systems are hard to get published in the scientific community. This policy clearly does not place incentives to undertake the effort of improving early prototype systems. As an editor of the KI journal, I cannot help but mention in this context that since about two years our journal features the article category of systems descriptions. The idea behind this article category is exactly to pave the path for systems that make the transition from scientific prototypes to mature tools. More visibility for such tools is important for their developers who get down to the gritty details of requirements from practice. There are no pearls without the grit.

Enjoy reading the current issue. Thank you!

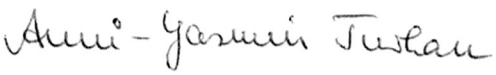

Anni-Yasmin Turhan

1 Call for Papers

1.1 Special Issue on Reintegrating Artificial Intelligence and Robotics

Guest Editors: Federico Pecora (Örebro University), Masoumeh Mansouri (Örebro University), Nick Hawes (University of Oxford), Lars Kunze (University of Oxford).

A major goal of artificial intelligence (AI) is to create autonomous, intelligent machines, or robots, that can sense their surroundings, reason about what they have perceived, plan their next actions, and act accordingly to accomplish their tasks. Moreover, robots should be able to learn from their own experience (including interactions with other agents) and adapt to changing conditions within their environments over their lifetime. It is the aim of this special issue to emphasize that the reintegration of AI methods is a non-trivial factor in the design, development and evaluation of robot systems. In particular, we are interested in work related to both fully-integrated robots systems that use methods of AI to perform complex tasks in realistic environments and fundamental AI techniques that have the potential to transform the capabilities of robot systems, but which not been convincingly demonstrated in integrated systems.

Topics of interest include, but are not limited to:

- Knowledge representation and reasoning for robots
- Qualitative representations for robots
- Integrated task and motion planning
- Context-based scene understanding
- Semantic mapping and reasoning with semantic maps
- Constraint-based reasoning for robots
- Continuous planning and on-line problem solving for robots
- AI-enabled human–robot interaction
- Lifelong learning and adaptation
- Verification of autonomous systems
- Reasoning with uncertain and inconsistent knowledge
- System-level AI for robots

If you are interested in contributing to this special issue, please contact one of the guest editors.

1.2 Special Issue on Artificial Intelligence in Games

Guest Editors: Sebastian Risi (IT University of Copenhagen) and Mike Preuss (University of Münster).

This special issue focuses on Artificial Intelligence (AI) methods applied in and for different types of games (e.g., board games, video games, serious games). Games have been shown to be the perfect testbed for advanced AI methods. AI in games is now a well established research area with two dedicated conferences and as well as a dedicated journal. Especially deep learning methods have recently proven to beat the best human experts in Atari video games and the game Go. Other methods such as evolutionary computation have been shown to allow complete new types of games through procedural content generation. While there has been much progress in game AI recently, some games such as StarCraft remain beyond even the most advanced AI algorithms. The goal of this special issue is to present a survey of the current research in Game AI and emerging trends in this area.

Some key topics of interest include:

- Procedural content generation
- Games as AI testbeds
- AI for NPCs (non-player characters)
- Player modelling
- Game mining
- Self-play
- Game balancing
- Human–computer interaction
- Human-based computation in games

If you are interested in contributing to this special issue, please contact one of the guest editors.

1.3 Special Issue on Challenges in Interactive Machine Learning

Guest Editors: Stefano Teso (KU Leuven, Belgium) and Oliver Hinz (Goethe University Frankfurt, Germany)

Machine learning is revolutionizing our ability to leverage data and tackle challenging applications such as natural language and image understanding, recommender systems, fraud detection, and medical diagnosis. Interactive learning supports these advances by casting learning as a dialogue between a model and one or more users, who may play the role of teachers, targets, and judges of the model being learned, and who can also learn as being part of the teaching–learning loop.

Designing successful interactive learning schemes requires to solve a number of key challenges:

- minimizing the cognitive cost for the user while optimizing query informativeness,
- devising effective interaction protocols based on different types of queries (membership, ranking, search, explanation, etc.),
- producing optimal questions by explicitly and efficiently capturing the uncertainty of the model,
- distributing the load of query answering across multiple teachers with heterogeneous abilities,
- designing or estimating realistic models of user behavior, increasing tolerance to noise and actively guiding the user toward providing better and more robust supervision, and, more generally,
- automatically discovering the user's expertise level and adapting the interaction accordingly. Such an interaction is likely to help make such systems more transparent and the results more explainable.

This special issue aims at surveying established research in interactive learning, as well as recent advances on algorithms, models and effective process design around humans in the loop. If you are interested in contributing (technical reports, system descriptions, project reports, survey articles, discussions and dissertation abstracts) to this special issue, please contact one of the guest editors before the submission deadline.