

2nd CALL FOR PAPERS

1st TAILOR Workshop at ECAI 2020, 4-5 September, Santiago de Compostela

Conference website: <https://liu.se/en/research/tailor>

Submission link: <https://easychair.org/conferences/?conf=tailor2020>

Submission deadline: June 4, 2020 (Anywhere in the world)

Notification of acceptance: June 23, 2020 (Anywhere in the world)

Camera ready paper: July 23, 2020 (Anywhere in the world)

The current scientific landscape is fragmented with many research groups working individually or in smaller constellations in often relatively isolated scientific communities: machine reasoning, machine learning, and optimization are examples of such mostly-disjoint communities.

The **TAILOR workshop** will bring these groups and researchers together in a unique atmosphere to discuss the state of the art and the **latest advances** in the **integration of learning, optimisation and reasoning** to provide the **scientific foundations for Trustworthy AI**. The TAILOR community builds upon a large H2020 proposal on building a network of centers of excellence that includes over 100 labs in Europe.

The **TAILOR network** of research excellence centers focus on the scientific foundations of **Trustworthy AI** integrating **learning, optimisation and reasoning**.

The main scientific topics of the workshop are:

- **Trustworthiness:** How to learn fair AI models, even in spite of biased data? How to develop explainable and interpretable AI decision processes? How to develop transparent AI systems and integrate them into the decision process for increasing user trust?
- **Learning, reasoning and optimisation:** How to integrate AI paradigms and representations for reasoning, learning and optimisation in order to support trustworthy AI? Integrated approaches to learning, reasoning and optimisation should allow AI systems to bridge the gap between low-level perception and high-level reasoning, to combine knowledge-based and data-driven methods, to explain their behaviour and allow for introspection of the resulting models.
- **Deciding and Learning How to Act.** How to empower an AI system with the ability of deliberating how to act in the world, reasoning on the effects of its actions, learning from past experiences, as well as monitoring the actual or simulated outcome of its actions, learning from possibly unexpected outcomes, and again reasoning and learning how to deal with such new outcome?
- **Reasoning and Learning in Social Contexts.** Agents should not reason, learn and act in isolation. They will need to do it with others and among others. So, this topic is concerned with how AI systems should communicate, collaborate, negotiate and reach agreements with other AI and (eventually) human agents within a multi-agent system (MAS).
- **AutoAI:** How to build AI tools, systems, and infrastructure that are performant, robust and trustworthy including having the ability to configure and tune itself for optimal performance? How can we support (1) people with limited AI expertise and (2) highly-skilled experts in building such AI systems?

We welcome submissions to the workshop of the following types:

1. Presentations of **relevant work that has recently been published or has already been accepted for publication** in journals such as AIJ, JAIR, JMLR, MLJ, and major conferences such as AAI, ICML, IJCAI, NeurIPS, SIGKDD, etc. The submission should in this case only consist of **a copy of the accepted paper**.
2. **Long papers reporting on new material**. Papers can be **at most 16 pages** in the Springer LNCS format. Please note that also shorter papers are welcome.
3. **Extended abstracts that report on novel and preliminary ideas**. Extended abstracts can be **at most 6 pages** in LNCS format.
4. **Short position statements on the topic of the workshop, at most 6 pages** in LNCS format.

The workshop is interested in foundational as well as more applied contributions. What matters is that they address the topics of the workshop - the integration of learning, optimisation and reasoning to provide the scientific foundations for Trustworthy AI.

Both long and short papers must be formatted according to LNCS guidelines (<https://www.springer.com/gp/computer-science/lncs/conference-proceedings-guidelines>) **and submitted electronically through EasyChair:** <https://easychair.org/conferences/?conf=tailor2020>

Publication

Proceedings: The papers accepted for TAILOR 2020 will be published in a Springer LNCS post proceedings.

Organisers

1. Workshop Chair Fredrik Heintz, fredrik.heintz@liu.se, Linköping University, Sweden
2. Luc De Raedt, KU Leuven, Belgium
3. Peter Flach, University of Bristol, UK
4. Hector Geffner, ICREA and Universitat Pompeu Fabra, Spain
5. Fosca Gianotti, National Research Council, Pisa, Italy
6. Holger Hoos, Leiden University, The Netherlands
7. Michela Milano, University of Bologna, Italy
8. Barry O'Sullivan, University College Cork, Ireland
9. Ana Paiva, University of Lisbon, Portugal
10. Marc Schoenauer, INRIA, France
11. Philipp Slusallek, DFKI, Germany
12. Joaquin Vanschoren, Eindhoven University of Technology, The Netherlands