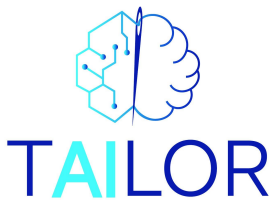

Foundations of Trustworthy AI – Integrating Reasoning, Learning and Optimization
TAILOR
Grant Agreement Number 952215

Third TAILOR Conference Report

Document type (nature)	Report
Deliverable No	1.5
Work package number(s)	WP1
Due Date	2023-08-31
Responsible Beneficiary	LiU, ID #1
Author(s)	Trine Vikinge, Fredrik Heintz
Publicity level	Public
Short description	The third out of four planned TAILOR conferences during the project

History			
Revision	**Date**	**Modification**	**Author**
1.0	2023-10-02	First version	Trine Vikinge, Fredrik Heintz

Document Review		
Reviewer	**Partner ID / Acronym**	**Date of report approval**
Janina Hoppstaedter	DFKI	2023-10-16
Luc De Raedt	KULeuven	2023-10-17

Table of Contents

Summary of the report

The third TAILOR conference was held June 5-6, in Siena, Italy. About 80 people participated, of which the majority were from project partners. The conference included thematic sessions on central topics like the AI Handbook, the roadmap and challenges, two poster sessions and two keynote presentations by external guests The conference setting provided plenty of space and opportunities for discussions and networking. The 6th project general assembly was held on June 5.

About TAILOR Conferences

TAILOR conferences are planned to be organised yearly. Partly, the purpose is to collect, integrate and share the scientific results and industrial applications. Another purpose is to update the roadmap by integrating the results from the research themes, other roadmaps and major initiatives, together with the needs from industry into a coherent and compact form guiding further research in the five themes and industry use cases. Furthermore, the TAILOR Conferences also constitute a way for partners to meet and network, being a space for discussion and scientific interaction.

Three TAILOR conferences were initially planned for the project. A fourth conference was added with the extension of the project from three to four years.
The 3rd conference was originally assigned to UNIROMA. In early discussions with the UNIROMA representatives, it became clear that Rome or the surroundings would not be optimal for the event. However, it was desired to keep the conference in Italy.

Representatives of the partner CINI, located in Siena, volunteered to take on the conference, offering to welcome us to the university conference venue; La Certosa di Pontignano. A compromise was made between easy access and atmosphere, as will be clear from the photos included in this report.

Conference	When	Responsible partner	Organiser
#1	M12	LiU	ULEI
#2	M24	LiU	CUNI
#3	M36	LiU	CINI
#4	M46	LiU	TBD

Conference preparations

Program committee

A conference program committee was established in February 2023 with the following members:

Name	Organisation
Fredrik Heintz (chair)	#1 LiU, leader of WP1 and in responsible for project conferences
Marco Gori	#37 CINI, head of local organisation
Peter Flach	#16 UNIBRIS, leader of WP9 Networking
Janina Hoppstaedter	#26 DFKI, substitute leader of WP8 Industry collaborations
Chokri Mrahda	#22 CEA, representing industry and involved in ADRA standardisation activities
Nikolaos Matragas	#22 CEA, representing industry and involved in ADRA standardisation activities
Francesco Giannini	#37 CINI, local organisation
Trine Vikinge	Project manager, #1, LiU
Lena Tasse	Communications director, #1, LiU

Planning of practical arrangements were started in fall 2022 in meetings between CINI, UNIROMA and LiU. The decision to move the assignment of the conference from UNIROMA

to CINI was included in the GA amendment filed in November 2022. During spring 2023, a number of meetings were held by the PC and selected members of it, depending on the matters to be discussed.

The format of the conference was copied from the TAILOR conference #2 held in Prague 2022, with some improvements based on the feedback received.

The program in the end covered the full days of the 5th and 6th of June, plus satellite meetings in the afternoon of the 4th.
VISION arranged the 3rd ICT-48 community workshop on June 7th, in Siena. Participants of the workshop were invited to join the conference in the afternoon of 6th June, as well as the social dinner closing the conference, and the program was planned taking this into consideration.

Conference venue

The conference venue was the lovely Certosa di Pontignano a few kilometres outside of Siena. (https://www.lacertosadipontignano.com/en/?r). All sessions as well as the welcome reception on the evening of June 4th were held here. Most participants stayed at the Certosa. However, around 15 people chose to stay in hotels in Siena town and bus transfer was arranged in the morning and evening.

The Certosa di Pontignano outside of Siena.

The CINI staff managed all local arrangements in close dialogue with the LiU project office.

Conference purpose

A key concept of TAILOR is the academic-public-industrial research network of the partners, with a combined capacity of providing the scientific basis for Trustworthy AI tailored to achieve both research excellence and broad participation with active industry involvement as well as providing active dissemination, education and training. A TAILOR conference is organised each year of the project to bring the community together to work on these matters. In general, the conferences are arranged to enable the collection, integration and sharing of scientific results and industrial applications. In order to enhance the networking and collaboration WP9 designed an activity to provide to each attendant a list of most similar attendees and their research work (additional details can be found in Appendix 3).

One of the main purposes of the conferences is to gather around the SRIR and plan for the next actions. A particular purpose of this year's conference was to shed light on AI regulation and standards development and to spur interest and engagement in this. Ray Walshe (Assistant Professor of ICT Standardization at Dublin City University, Ireland) was invited and confirmed but cancelled on short notice for health reasons. Francois Terrier, VP for AI program and Director of CEA List, was invited to replace Ray and gave a much appreciated remote keynote on the topic.

The conference was in-person, except for the two keynotes, a brief meeting with the EC PO, and the introductory presentation by project coordinator Fredrik Heintz. The introduction was broadcasted on zoom and open to all, but it was particularly meant for the PIs that could not be present in person for the General Assembly, to keep the Assembly meeting shorter.

Poster sessions were included in the program, on both days. All conference participants were welcomed to bring posters and particularly young researchers and industry representatives were encouraged. Re-use of previously exposed posters were welcomed and no selection process was applied.

Conference audience

The conference is the main opportunity offered for partners and network members to meet in person and interact. Messages were issued from the project coordinator to underline the importance of participating; that being a partner in the project obliges them to participate in these events. The ambition was to have 2-3 people per partner and 1 person from each of the ca 90 network members, aiming for ca 250 participants.

The conference was extensively announced, starting from February 2023, using the TAILOR web-page, social media, the TAILOR newsletter, and group mail to the target groups. Ca 85 people had registered to participate a week before the event. In the end, 80 people participated, from 35 project partners, including 4 industry partners, and from 1 industry network member.

Registration for the conference was opened in March and closed at the end of May.

Registration of participation, posters, and conference dinner were managed by an online tool by LiU project office.

Pre-Conference Program

On Sunday June 4th, two workshops were arranged; one joint for WP4 Paradigms and WP5 Acting, and one for WP7 AutoAI. Altogether, around 25-30 participated in these.
A welcome reception and dinner was offered to all who arrived on Sunday June 4.

WP4 and WP5 combined workshop

This workshop had a goal to foster interaction between the two WPs 4 and 5, with a special interest on how to use integrated representations for acting.

The workshop started with a keynote presentation by Hector Geffner who set the scene (in an online presentation). Agnieszka Lawrynowicz presented some work connecting to natural language processing. In the discussion that followed, ideas for potential collaboration were shared and the influence of large language models on the research was discussed.

WP7 AutoAI workshop

This in-person WP7 meeting marked an invaluable opportunity for task leaders to share the latest scientific advancements emerging from their respective research groups. The intimate scale of the workshop fostered an environment where discussions thrived. Ideas were exchanged and insights shared.
After the presentations were finished, a walking discussion on next steps for the work package was arranged.

.

Conference program

The conference schedule for 5th and 6th June is shown in Appendix 1.

The program started both days at 9am and ended at 6pm. On June 5th, the General Assembly was held from 6pm to 6.30pm.

Coffee was served in the morning and in the afternoon and lunch was served mid-day. Dinner was served at 8pm.

Coffee, lunch and dinner were complementary and on-site, which assured that the discussions in the conference auditorium continued into and throughout the breaks. Also, posters were exposed in the area where coffee and meals were served, which assured constant gathering of people around these.

Day 1, June 5

Welcome and TAILOR update

The coordinator Fredrik Heintz opened the meeting and welcomed all. He gave a presentation of the project; refreshing our memories on the project goals and objectives, and summarising the achievements and technical highlights so far. A few questions for further discussion during the conference concluded the presentation:

1. How to show and quantify the impact of TAILOR?
2. What long-term role could TAILOR play? How can we sustain the network after the project ends?
3. What are the most important things to focus on now, with one year left of the project?

The introductory presentation was made available on zoom for all partners not participating on site and all general assembly representatives not on site had been encourage to watch, so the same presentation would not have to be given at the general assembly #6 later the same day.

Poster sessions 1 and 2

There were two poster sessions on day 1; one in the morning and one in the afternoon. The first poster session started with lightning pitches in the auditorium; all poster presenters were given a 60 seconds slot, with or without a slide, to tease the audience.
33 posters were presented and the presenters names and poster titles are shown in Appendix 2. A poster booklet from the conference can be found on the TAILOR website.
The pitches were followed by ample time by the posters.

Session on the TAILOR Handbook of Trustworthy AI

The TAILOR Handbook of Trustworthy AI is an encyclopaedia of the major terms related to AI trustworthiness. The Handbook was a deliverable at M18 (D3.3) and is also publically accessible as a WIKI. A second version is due M46. Partner #2 CNR, led by Umberto Straccia and Francesca Pratesi, is in charge of WP3 and were responsible for this session at the conference

In this session, all participants were assigned to a group, and each group was assigned one of the specific ethical dimensions of the Handbook: Explainable AI systems, Safety and robustness, Fairness, equity, and justice by design, Accountability and reproducibility, Respect for privacy, Sustainability.
The groups were instructed to list five or more fundamental properties characterising such dimensions, then look for it in the Handbook and discuss, and towards the end of the session return with suggestions for additions, modifications, and improvements.

The group work was collected by the WP3 management team and will be used by the Handbook task force in the preparation of the version 2 of the Handbook..

Challenges

An objective of WP2 is to guide and boost research and innovation by setting up benchmarks and competitions. These are intended to contribute to tackling the trustworthiness issues, as well as to pushing toward the integration of learning, reasoning and optimization.
Marc Schoenauer, Wico Mulder and Marco Gori held this session, presenting what has been done and what is coming up.
The project plan is to run, analyse and evaluate at least 3 academic and 3 industrial benchmarks.

Challenges successfully completed within TAILOR were presented:

- Smarter Mobility Data Challenge, EDF + Manifest IA (Oct. Dec. 2022)
  - Final jury in March 2023, pdt Cédric Villani
- L2RPN II: Toward Carbon Neutrality, RTE (June-Sept. 2022)
  - and the winner is … a mix of Expert Rules and Deep RL
- MetaLearn 2022, Inria partner et al., TAILOR-sponsored prizes (Summer 2022)
  - Learning from Learning Curves - perf. w.r.t. dataset size (AutoML)
  - Cross-domain MetaDL - Any way/any shot meta learning
- Brain Age Prediction from EEG Challenge, NeuroTechX (Nov. 2022)
  - discrepancy with actual age is a critical psychiatric indicator

Some challenges coming soon were also presented:

- Mind your buildings
  This challenge is about identifying behavioural patterns related to building occupancy using sensor data coming from a multi tenant building. In the period from January to March 2023, a group of 25 people worked on data science problems in the context of urban energy sustainability. It was organised by TNO and DFKI, in collaboration with the Hanze university of applied sciences in the Netherlands and the company AIMZ.

The groups developed algorithms that could pinpoint and repair missing data in incomplete floor plans of buildings. Also, models for prediction of occupancy were retrieved from the sensor data.

During the meeting in Siena, it was explained that the topic of Theory of Mind was a challenging topic to tackle during this hackathon, therefore a follow up challenge was announced, intended to be held under the title *Mind the avatars' mind*.

- Crossword puzzle
  Organised by Prof. Marco Gori's WebCrow team at U. of Siena, this challenge has two phases, addressing automated crossword solving and generation, based on common modules hybridising Natural Language Understanding (NLU), Machine Learning and constraint satisfaction, while gathering knowledge and data from several sources (web search, dictionaries, specialised multilingual schools curricula). Understanding crossword definition goes beyond NLU: Understanding clues requires several logical steps in Language Analysis.
  Crossword solving will involve gradual tasks, from traditional clue answering and grid filling to integrated approaches for constrained clue answering, crossword correction, and end-to-end Neuro-Symbolic models.
  Crossword generation will be personalised (finding topic-relevant terms and clues/definitions), and will involve the design (or fine-tuning) of some Vertical LLM for direct generation of clue/answers.

- ML for Physical Simulations (aka Scientific Machine Learning – SciML)
  Co-organised by TAILOR (through its Inria partner) and several industrial partners (including NVIDIA, RTE and Criteo), this challenge intends to promote the use of Machine Learning based surrogate models to numerically solve physical problems, through a task addressing a Computational Fluid Dynamics (CFD) use case related to airfoil modelling. The challenge will be held on the Codalab platform (maintained by the Inria partner), from Nov. 16. 2023 to end February 2024. The public training dataset is the AirFrans dataset described in the NeurIPS (dataset and benchmarks track) paper, made of 1000 CFD simulations of steady-state aerodynamics over two dimensions airfoils in a subsonic flight regime (5 real values at every point of the point cloud defined by the mesh on the simulation domain), and the participants will have access

for their simulations to the LIPS (Learning Industrial Physical Simulation) platform described in the NeurIPS (dataset and benchmarks track) paper. The task will be to build surrogate models of these 5 fields for new airfoils, including Out-ot-Distribution cases. The evaluation will be a mix of accuracy (MSE), computational cost, and, last but not least, respect of the physical constraints (Navier-Stokes equations).

A few challenging aspects were discussed, like lack of volunteers / time / energy
    - In particular for Hackathons …
- Links with TAI through reproducibility and fair results assessment
- LOR combination
    - not enforced
    - but present in most winning solutions, see forthcoming deliverables

General assembly #6

The General Assembly is arranged every six months and this one was the first one in-person. Altogether, 39 partners were present, the majority on-site and a minority on zoom. Minutes from the General Assembly are confidential and not included in this report.

Conference dinner

A very memorable dinner was arranged for all conference participants at the Certosa in the evening of June 5.

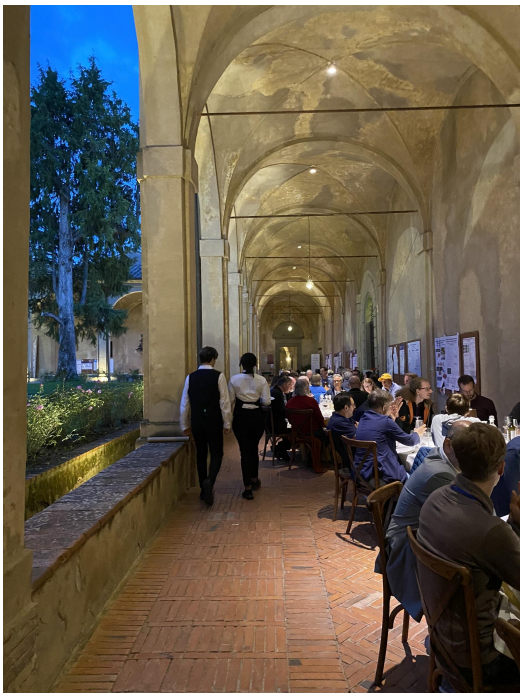
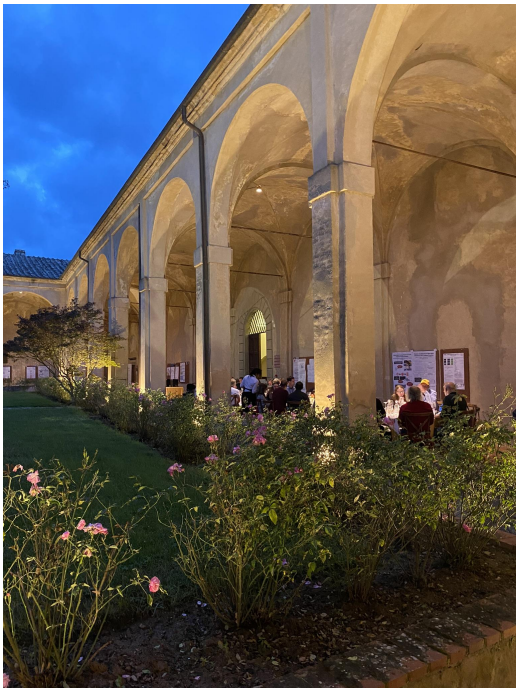

Day 2, June 6

The second day of the conference was dedicated to industry, roadmaps and looking ahead. The morning was dedicated to industry and particularly to the importance of regulation and

standards and their implications. The afternoon had a focus on what lies ahead, related to the SRIR version 2 and the role of TAILOR with regard to large language models. The afternoon was designed to be of interest to the other NoEs, whose partners would be in the area for the ICT-48 workshop the day after.

Industry session

The leader of WP8, Industry, Innovation and Transfer program, Philipp Slusallek, introduced the morning industry session with the theme AI regulation and standards and their importance for technology transfer. First on the agenda was a keynote presentation by Francois Terrier, CEA. Thereafter, four panellists each gave a 15 minutes presentation on the topic, before a panel discussion in which also the audience was involved.

Keynote by Francois Terrier: Regulation, Standards, Methods and Tools for AI Safety

Francois Terrier stepped in on short notice in place of Ray Walshe, who had to cancel for private reasons. Francois Terrier is the VP for AI program and Director of CEA List at the TAILOR industry partner CEA

Presentation title: No Trust without regulation! European challenge on regulation, liability and standards for trusted AI.
Abstract: The explosion in the performance of Machine Learning (ML) and the potential of its applications are strongly encouraging us to consider its use in industrial systems, including for critical functions such as decision-making in autonomous systems. While the AI community is well aware of the need to ensure the trustworthiness of AI-based applications, it is still leaving too much to one side the issue of safety and its corollary, regulation and standards, without which it is not possible to certify any level of safety, whether the systems are slightly or very critical.
The process of developing and qualifying safety-critical software and systems in regulated industries such as aerospace, nuclear power stations, railways or automotive industry has long been well rationalised and mastered. They use well-defined standards, regulatory frameworks and processes, as well as formal techniques to assess and demonstrate the quality and safety of the systems and software they develop. However, the low level of formalisation of specifications and the uncertainties and opacity of machine learning-based components make it difficult to validate and verify them using most traditional critical systems engineering methods. This raises the question of qualification standards, and therefore of regulations adapted to AI. With the AI Act, the European Commission has laid the foundations for moving forward and building solid approaches to the integration of AI-based applications that are safe, trustworthy and respect European ethical values. The question then becomes “How can we rise to the challenge of certification and propose methods and tools for trusted artificial intelligence?”

> **Bio**: François Terrier is AI Senior Fellow at CEA. He has a PhD in artificial intelligence and worked 10 years in the domain of expert systems using three-valued, temporal or fuzzy logics. As head of the system and software engineering department of the LIST Institute of CEA Tech since 1994, he led actions to build, for trustworthy software and systems, open tool chains covering the whole development

cycle from requirement specification until equipment integration. His research challenges are, namely, on combining domain oriented modelling with formal methods for high quality, safe and secure critical systems. Since 2020, François is CEA VP for Artificial Intelligence program and the Director of the programs at CEA institute for smart digital systems (CEA List).

Industry panel

The industry panel had four participants, listed in the table below.

Organisation	Name	Presentation title
CBS	Barteld Braaksma	AI, Official Statistics and Standards
Tieto	Ifthikar Ahmad	Building trusted digital societies
#39, TNO	Wico Mulder	Three pillars to bridge the GAP to industry
#1 Liu	Fredrik Heintz	Adra and the EU AI, Data and Robotics Public-Private Partnership

The discussions that followed the panellists' presentations were vivid and also engaged the audience. Some more or less concrete suggestions were put forward:

1. The project should make efforts to create a push for industry transfer in the last year. To achieve this, it was suggested to
    a. redistribute some of the project funding for strengthening the incentives for collaboration between industry and academy.
    b. Set up joint projects and push for internships with industry
    c. Establish bases for longer term collaboration, such as within the frames of so-called Transfer Labs:
    d. Enable bilateral collaboration on industry data
2. On the matter of standardisation efforts, it was proposed to engage through specific, time-boxed task forces
    a. An example of such a time.boxed initiative is the WP of ADRA-e concerned with this.
3. On the topic of interactions with industry, it was discussed how to make the competencies and capacities of the academic world available to industry. Ideas that came up were for example
    a. Webpage offering our high impact research results, tools, data sets, success stories, etc.
    b. To create interface to industry with a contact team in TAILOR for answering and distributing requests
    c. Matchmaking between industry and interested researchers
    d. To actively involve applied research centres (CIIRC, DFKI, FBK, FhG, Inria, TNO)
    e. Promote this widely as a central resource for European industry
    f. Offer to include other networks

g. It was stated that sustainable and long-term value demands collaboration with permanent initiatives such as CLAIRE.

Already during the conference first ideas for the implementation of the above mentioned points were discussed. Based on this, funding opportunities and possible collaborations are currently being explored. Further meetings and sightings of existing solutions will be sought before the end of 2023.

Latest news from DG CNECT; Meeting with the EC PO

The Project Officer of the EC, Evangelia Markidou, honoured us with her presence on zoom for a 30 minutes discussion in the afternoon of Day 2. The time slot was dedicated to forward looking discussions and the view of the Commission on Trustworthy AI, particularly in relation to large language models..

Road map and joint SRA

A Strategic Research and Innovation Roadmap (SRIR) session was organised and led by Michela Milano and Roberta Calegari, #10 UNIBO, who are in charge of gathering input for the SRIR v.2 through interactions with other initiatives.

The TAILOR Strategic Research and Innovation Roadmap (SRIR) aims to boost research on Trustworthy AI by clearly defining the major research challenges. The first version was completed spring 2022 and the second version will be completed spring 2024. Additionally, together with the other initiatives (the 4+2 NoEs and VISION[1]), a Joint Strategic Research Agenda (SRA) is produced. The first version was completed in June 2023 and the second

[1] The Coordination and Support Action VISION: https://www.vision4ai.eu/
The original 4 NoEs: AI4Media: https://www.ai4media.eu/ , ELISE: https://www.elise-ai.eu/ HumanAINet: https://www.humane-ai.eu/
Two new NoEs: euRobin: https://www.eurobin-project.eu/ , ELSA: https://www.elsa-ai.eu/

version will be completed June 2024. The Joint SRA is broader than the SRIR, reflecting the breadth of the 6 NoEs.

The session addressed pressing questions concerning the measurement and quantification of Trustworthy AI (**TAI**), trustworthiness certifications, and the necessary mentoring and training required in the context. Specific questions were devoted to delving into the interconnectedness of Learning Optimization Reasoning (**LOR**) and TAI, examining the pathways from LOR to TAI and vice versa.

Topics discussed:

1) Measuring and Quantifying Trustworthy AI: One of the central topics of the session was the challenge of measuring and quantifying TAI. Participants acknowledged that ensuring trustworthiness in AI systems is a multifaceted endeavour. Key points of discussion included the need for standardised metrics and evaluation frameworks to assess the trustworthiness of AI algorithms and systems comprehensively. This would require a combination of technical assessments, ethical considerations, and real-world impact assessments.

2) Trustworthiness and Certifications: The session highlighted the increasing demand for trustworthiness certifications in the AI industry. Participants discussed the importance of establishing industry-wide certification standards that AI developers and organisations must adhere to. These certifications would serve as a mark of trust for consumers and regulators, helping to bridge the gap between technical excellence and societal acceptance of AI technologies.

3) Mentoring and Training: Recognizing the rapid evolution of AI technologies, the session emphasised the critical need for mentoring and training programs. It was noted that the AI community should invest in mentorship initiatives to guide new talent and foster ethical practices. Ethical considerations and responsible AI development should be core components of AI education and training programs to ensure the next generation of AI professionals is well-equipped to handle the challenges of trustworthiness.

A specific session was devoted to the interconnection between LOR and TAI: Participants explored the intricate relationship between Learning Optimization Reasoning (LOR) and Trustworthy AI (TAI) from two perspectives:

4) From LOR to TAI: The discussion revolved around how to make LOR-based systems trustworthy. It was emphasised that integrating ethical principles, bias mitigation techniques, and robustness measures into LOR algorithms is crucial. Ensuring transparency, interpretability, and accountability in LOR-based AI models emerged as key priorities.

5) TAI via LOR: The reverse pathway, enforcing TAI in AI systems that exploit LOR, was also examined. Participants stressed the importance of using LOR techniques to enhance the trustworthiness of AI systems. By employing LOR to optimise decision-making, reasoning, and ethical considerations, AI systems can become more reliable, transparent, and aligned with human values.

In conclusion, the SRIR session tackled critical questions surrounding TAI, trustworthiness certifications, mentoring, and the intersection of LOR and TAI. The consolidation of these insights into the final version of the SRIR will guide the AI community toward a future where AI systems are not only technically proficient but also trustworthy, ethical, and aligned with societal values.

Role of TAILOR in respect of Generative AI

Panel discussion moderated by Fredrik Heintz; The purpose of the panel is to discuss the role of TAILOR with respect to Generative AI in Europe. We will cover matters such as: What are the main challenges for Generative AI in Europe? and What are the opportunities for TAILOR wrt to Generative AI? The conclusion is that TAILOR has a lot to offer as today's generative AI methods and models are far from trustworthy. Most likely it will also be necessary to take advantage of techniques from reasoning to complement the machine learning methods used.

Keynote by Thomas G. Dietterich: Competence Models For Machine Learning Systems

The conference was ended by a keynote speech by Thomas G. Dietterich, distinguished Professor Emeritus, competence models for machine learning systems.

Keynote abstract: Every AI system, and certainly every machine learning system, should have an accurate model of its own competence. This talk will review our recent work on creating competence models for classifiers, computer vision systems, and reinforcement learning. We will discuss both input competence–is the system competent to handle a given query–and output competence–can the system provide calibrated statements about the quality of the answer it has produced? For input competence, our focus is on detecting queries that fall outside of the training data. We apply anomaly detection algorithms to address this task. For manually-engineered features, these work well, but we have discovered that features learned by deep neural networks are often inadequate for this task. The problem is that features are only learned for "directions of variation" that are present in the training data. An anomaly that varies in some way that was invariant in the training data will not be detected. For output competence, we have studied calibration (for classifiers) and trajectory-wise prediction intervals (for reinforcement learning). We will report progress on point-wise calibration, which is more challenging that the standard notion of set-wise calibration. Then we will describe our approach to creating trajectory-wise prediction intervals using the tools of conformal prediction. All of these results assume stationarity of the distribution that is generating the data. We will conclude by briefly discussing our efforts to extend this work to shifting distributions.

> **Bio**: Dr. Dietterich (AB Oberlin College 1977; MS University of Illinois 1979; PhD Stanford University 1984) is Distinguished Professor Emeritus in the School of Electrical Engineering and Computer Science at Oregon State University. Dietterich is one of the pioneers of the field of Machine Learning and has authored more than 200 refereed publications and two books. His current research topics include robust artificial intelligence, robust human-AI systems, and applications in sustainability.

Dietterich has devoted many years of service to the research community and was recently given the ACML Distinguished Contribution and the AAAI Distinguished Service awards. He is a former President of the Association for the Advancement of Artificial Intelligence and the founding president of the International Machine Learning Society. Other major roles include Executive Editor of the journal Machine Learning, co-founder of the Journal for Machine Learning Research, and program chair of AAAI 1990 and NIPS 2000. He currently serves as lead moderator for the cs.LG category on arXiv.

Conference Closing

Fredrik Heintz closed the meeting, thanking all, and welcoming to the social dinner.

Social dinner

All participants of the TAILOR conference as well as participants of the ICT-48 workshop the day after were invited to a social dinner in the fantastic Tuscan countryside. Buses were chartered for transport from La Certosa to Brolio, located in the midst of the beautiful Chianti landscape. A walk around the Castello di Brolio was organised before dinner at the Osteria di Brolio, which is part of the Ricasoli wine producing establishment. Transport back to both La Certosa and Siena after dinner was arranged. Ca 80 people participated at dinner, of which ca 10 were from other NoEs or VISION.

Appendix 1: Program

Day 1, June 5

Time	Title	Speaker
08:00	Registration	
09:00	Welcome and introduction	Fredrik Heintz
09:15	Slot 1 - TAILOR - What have we achieved? Where are we going?	Fredrik Heintz
10:00	Coffee	
10:30	Slot 2 - Poster session 1 (first half presents)	
12:30	Lunch	
14:00	Slot 3 - Handbook on Trustworthy AI / Workshop on Wikipedia	Lena Tasse
15:30	Coffee	
16:00	Slot 4 - Poster session 2 (other half presents)	
17:45	Challenges presentations	Marco Gori & Marc Schoenauer
18:00	General Assembly	Fredrik Heintz
18:30	Break before Dinner	
20:00	Conference dinner at La Certosa	

Day 2, June 6

Time	Title	Speaker
09:00	Introduction Industry session	Philipp Slusallek, DFKI
09:05	Slot 5 -Francois Terrier; Regulation, Standards, Methods and Tools for AI Safety	
09:50	Slot 6 -Panelists introduce themselves	

10:30	Coffee	
11:00	Panel discussions, moderated by Philipp	
12:30	Lunch	
14:00	Slot 7 - EC presentation	Evangelia Markidou, EC
14:15	Slot 8 - Road map and Joint SRA	Michela Milano, Fredrik Heintz
15:30	Coffee	
16:00	Slot 9 - Panel discussion: Role of TAILOR in respect of LLM	Moderator: Fredrik Heintz
17:00	Slot 10 - Keynote by Thomas Dietrich, TAILOR IAB-member.	
17:55	Closing	Fredrik Heintz
18:00	Break before dinner	
19:00	Dinner at a local winery	

Appendix 2: Posters

First name	Last name	Affiliation	Title
Annelot	Bosman	Universiteit Leiden	Preliminary work on robustness distributions in neural network verification
Barteld	Braaksma	Statistics Netherlands	AI and Official Statistics
Franco Alberto	Cardillo	Consiglio Nazionale delle Ricerche, Istituto di Linguistica Computazionale, Pisa, Italy	PN-OWL: A Two Stage Algorithm to Learn Fuzzy Concept Inclusions Rules from OWL Ontologies
Luciano	Cavalcante Siebert	TU Delft	Meaningful human control: actionable properties for AI system development
Francisco	Chicano	University of Malaga	Time Series Forecasting for Parking Occupancy
Francisco	Chicano	University of Malaga	Surrogates for Permutation Problems
Stefano	Fioravanti	CINI	Categorical Foundation of Explainable A
Francesco	Giannini	CINI - University of Siena	Relational Reasoning Networks
Ali	Kordia	GAIPS, Instituto Superior Técnico	Recognition and Prediction Using Dynamic Movement Primitives
Krzysztof	Krawiec	Poznan University of Technology	Self-supervised Learning of Tokenized Representations for Solving Raven Progressive Matrices Problems
Krzysztof	Krawiec	Poznan University of Technology	Learning to Solve Abstract Reasoning Tasks with Neurosymbolic Program Synthesis
Chokri	Mraidha	CEA	Towards a Framework for Uncertainty-Aware Dynamic Risk Management for Deep Learning-Based Autonomous Systems

Andrea	Orlandini	National Research Council of Italy	Optimal task and motion planning and execution for human-robot multi-agent systems in dynamic environments
Peter	Pavlík	slovak.AI	Physics-informed Machine Learning in Precipitation Nowcasting
Miquel	Perello-Nieto	University of Bristol	Classifier Calibration: A survey on how to assess and improve predicted class probabilities
Steven	Prestwich	UCC	Solving Complex Optimisation Problems by Machine Learning
Fabrizio	Riguzzi	CINI	New Benchmarks for Neuro-Symbolic Systems
Emanuele	Sansone	KU Leuven	Learning Symbolic Representations Through Joint GEnerative and DIscriminative Training
Jonathan	Shaki	Bar-Ilan University	Cognitive Effects in Large Language Models
Carles	Sierra	IIIA-CSIC	Enabling Game-Theoretical Analysis of Social Rules
Piotr	Skrzypczynski	Poznan University of Technology	Fast Motion Planning with Deep Neural Networks
JERZY	STEFANOWSKI	Poznan University of Technology	Multi-criteria approach for selecting an explanation from the set of counterfactuals
Stefano	Teso	University of Trento	Neuro Symbolic Continual Learning: Knowledge, Reasoning Shortcuts and Concept Rehearsal
Vito	Trianni	ISTC, Consiglio Nazionale delle Ricerche	Hybrid Collective Intelligence for Decision Support in Complex Open-Ended Domains
Veronique	Ventos	NukkAI	Explainability à la NukkAI
Andrea	Visentin	University College Cork	AI-based Task Classification with Pressure Insoles for Occupational Safety
Neil	Yorke-Smith	TU Delft	Adaptive Parallelization of Multi-Agent Simulations with Localized Dynamics
Neil	Yorke-Smith	TU Delft	Predicting the Optimal Period for Cyclic Hoist Scheduling Problems
Kamyar	Zeinalipour	University of Siena	Graph Neural Networks for ECG Classification]{Graph Neural Networks for Topological Feature Extraction in ECG Classification
Umberto	Straccia	CNR	Explainable Malware Detection

Robin	Manhaeve	KU Leuven	Advances I Neural-Symbolic AI
Daniela	Tulone	CSAIL, MIT	#D Tech-responsible Design to Boost AI Sustainability Benefits while Mitigating Risks
Marc	Schoenauer	INRIA	Memetic Semantic Genetic Programming for Symbolic Regression

Appendix 3: Networking and collaboration activity

The Work Package 9 Team at UNIBRIS (Miquel Perello-Nieto and Peter Flach) is committed to enhancing the collaboration and networking between TAILOR members, partners and others. Some of the main drivers of the team are the development and dissemination of AI-powered collaboration tools (Task 9.1) and the summer schools and other training events (Tasks 9.4). The WP9 team successfully used AI tools in previous Summer Schools and TAILOR Conferences. During the 2nd TAILOR Summer School (Barcelona, June 2022) WP9 used SubSift [1] and SynTeams [2] to match the interests and expertise of students and create synergetic groups that could lead to future collaborations. After the activity, 69% of the students that provided feedback indicated the consideration of using the TAILOR Connectivity Fund to visit a TAILOR lab (or host a student in their TAILOR lab). 75% of those with students in their assigned groups. The same tools were used during the 2nd TAILOR Conference (Prague, September 2022) to create synergetic groups of similar expertise to work on the Strategic Research and Innovation Roadmap. The feedback was overall positive with 96% of the responses indicating having expanded their professional network.

For the 3rd TAILOR Conference, WP9 used SubSift and AI explainability tools to discover potential connections and enhance discussions among attendees. The main idea was to characterise every participant with the entries in their computer science bibliography (personal DBLP webpage), or a free text description of their previous work and interests. Of 86 registered attendees to the conference, 40 agreed to participate in this activity through filling out a web form.

The WP9 team used a web crawler to automatically extract the list of entries in the DBLP webpage of each participant including titles of journal articles, conference and workshop papers, parts in books or collections, informal and other publications, reference works, editorships, books, theses among others. The team used SubSift to extract the most meaningful tokens of text to characterise the participants. The characterisation involved the computation of TF-IDF (Term Frequency - Inverse Document Frequency) and the standardisation of each participant feature representation to unit length (L2 normalisation). Then, each participant was scored with the rest of the participants using cosine similarity (see *Figure 1* for an example of the ranking).

In order to promote discussions about previous work and stimulate future collaborations, the team used SHapley Additive exPlanations (SHAP) [3] to highlight the entries in their DBLP matches that contributed positively (or negatively) to each score (see *Figure 2* for an example and extended explanation). The SHAP value indicates how the score changes when a particular piece of text is removed to a participant's profile.

On the 2nd of June (3 days before the conference) all participants were contacted and a unique URL was provided to each of them with an explanation of the activity (for transparency) and their ranked scores with other participants. During the conference, participants could access their results and explore the titles of other participants to facilitate discussions and potentially promote future collaborations.

During and after the conference the team collected feedback from the participants by informal discussions and a formal questionnaire to be filled by the 22nd of June. The feedback was overall positive with a particular interest in the used methods and the information obtained in their highest matches. The questionnaire was filled by 27.5% of the participants (11 out of 40) which provided the following insights. 72.7% of them checked the results before the conference (8 out of 11), and 18.2% after the conference. The provided matches to each participant were perceived as a correct match by 72.7% with the rest of them not having a strong opinion. 27.3% discovered unexpected related work with other researchers and another 27.3% partially agreed. 63.6% considered that the information provided may lead to future collaborations. Finally, all respondents considered that the information provided was somehow useful and that this type of activity may be useful in larger conferences where the majority of attendees are unknown. The following are two examples of feedback provided in the questionnaire:

"Smooth process and interesting experience."

"Interesting to meet people that work in related fields."

- Score 0.38 - Luc De Raedt
- Score 0.37 - Saso Dzeroski
- Score 0.30 - Miquel Perello-Nieto
- Score 0.30 - Joaquin Vanschoren
- Score 0.21 - Fabrizio Riguzzi
- Score 0.20 - Stefano Teso
- Score 0.20 - JERZY STEFANOWSKI
- Score 0.17 - Francesco Giannini
- Score 0.17 - Roberta Calegari
- Score 0.15 - Maria Bielikova

Fig 1: *Example of 10 participants with the highest similarity score based on the titles of publications from DBLP for the user Peter Flach.*

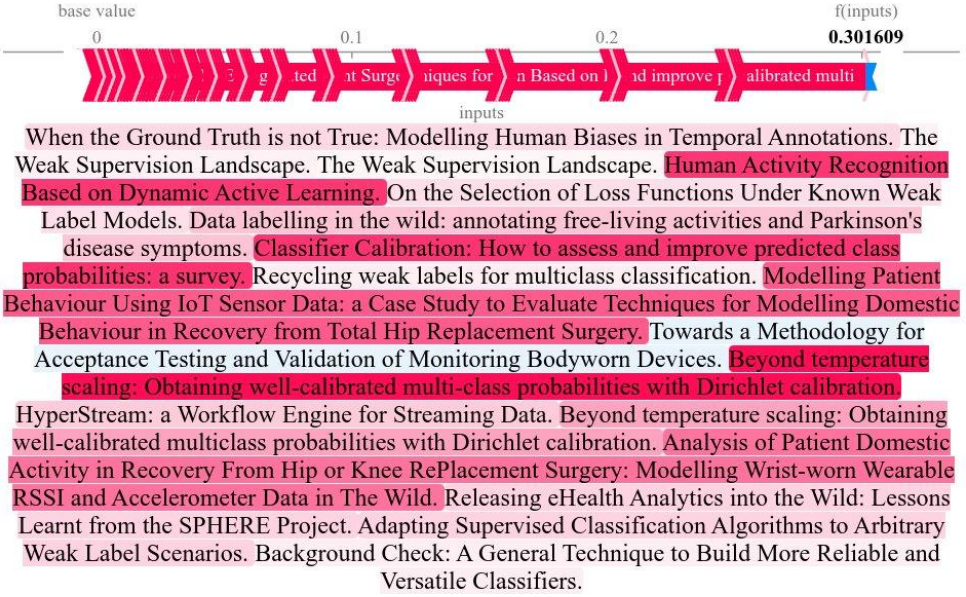

Fig. 2: *Example of explainability method SHAP with the DBLP profile of Miquel Perello Nieto to explain the similarity score of ~0.3 with Peter Flach. The top axis shows the base score value of zero on the left (average similarity score of Peter Flach and all the participants), and the similarity score of ~0.3 on the right side for Miquel. The highest positive contribution (red) comes from the conference paper "Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration" with a positive contribution of 0.054. While the lowest negative score (blue) of -0.006 is obtained for the paper "Towards a Methodology for Acceptance Testing and Validation of Monitoring Bodyworn Devices."*

References

[1] Price, S., Flach, P.A. & Spiegler, S.. (2010). SubSift: a novel application of the vector space model to support the academic research process. *Proceedings of the First Workshop on Applications of Pattern Analysis*, in *Proceedings of Machine Learning Research* 11:20-27.

[2] Ewa Andrejczuk, Juan A Rodríguez-Aguilar, Carles Sierra, Carme Roig, and Yolanda Parejo-Romero. Don't Leave Anyone Behind: Achieving Team Performance Through Diversity. In 2018 IEEE Frontiers in Education Conference (FIE), pages 1–9, 2018.

[3] Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc.