Let's Talk AI with Tiziana Margaria

Tiziana Margaria[1] and Barbara Steffen[2]

[1] University of Limerick, Department of Computer Science and Information Systems, Ireland,
[2] tiziana.margaria@ul.ie
[3] METAFrame Technologies GmbH,
barbara.steffen@metaframe.de

"I worry that improved AI will foster more delegation to it, where humans become lazier and just use "the machine", in delegation beyond what's necessary for a superficial sense of "convenience". That's a scary possibility for our future: a very likely future, I am afraid."

The Interviewee - Tiziana Margaria

My Personal AI Mission:
Do the possible so that AI is not taken as an "oracle", i.e., as an entity that does not need to be doublechecked or questioned. I'd prefer AI as Augmented Intelligence: considering intelligently what best to augment and how, and how to keep it in check and monitored, instead of (blind) Artificial Intelligence substituting the humans.

My Takes on AI

Artificial Intelligence: Right now, a set of deductive or associative systems able to respond to a series of inputs with adequate outputs, at least most of the time (to include nondeterministic and probabilistic systems like most AI techniques, which are wrong a significant number of times).

There are as of today no agreed thresholds when intelligence is/becomes "artificial", nor when an algorithm or a system is/becomes "intelligent". The Turing

test is in my opinion not exhaustive, as it is full of systems where a classic algorithm takes a decision as well as a human and is therefore indistinguishable (thus satisfying the Turing definition), but nobody would call that algorithm to be "artificial intelligence". Examples are a simple pocket calculator, a thermostat for a heating controller. If we exclude algorithms, it's difficult to save Neural Networks and other similar systems, which are de facto algorithms, just with weights that depend on the inputs and some feedback loops and stabilize after a while. The concept of "learning" is also not clearly defined for this use: if it is reproductive learning, like traditional training-based AI, it errs when extrapolating in new situations (in other words, it does not experience enough to be useful in general). If we go to the LLMs, as I heard most recently from Edward Lee at DATE 2024, their training is like educating by exposure to information, experience, and feedback. Is that artificial intelligence? Is education in general then "artificial"? In which sense is the process to train an LLM then similar or different to educating a child? I think, maybe new categories of natural/artificial and of intelligence will emerge in due time. We used to distinguish between small scale neural networks and deep learning, now LLMs seem definitely to be a different "thing" from both of them.

Trust: Confidence that the entity in question (a human, an algorithm, or a black-box like AI, the good old expert systems, or the Mechanical Turk of Amazon's history) does not deceive. I think this encompasses more than "tells the truth". It's a matter of having the right knowledge, applying it in toto, and not misrepresent or deviate.

Explainability: Ability to justify each individual recommendation on the basis of the inputs (that include the context) and some reasoning (deductive, abductive, or through some mapping of the input/output space, e.g., through decision trees).

Essential Elements of Human Capabilities: Conscience, feelings, moral. Associative abilities are now also shown by algorithms or agent societies like swarms, so these are not anymore peculiar to the humans. Intelligence is also debatable: the ability to predict (mostly by extrapolation from experience) is now seemingly the "big thing" of AI too: given experience with training data, make prediction on fresh data. Human intelligence would usually also use other criteria like what's legal, what's opportune etc. which go beyond just looking at sets of data. With the LLMs, however, so much input is provided upfront during the training that such side criteria might well be meanwhile included in the answers.

The Interview

Barbara *Today I have the pleasure to interview Professor Tiziana Margaria from the University of Limerick. Please briefly introduce yourself and your relationship to artificial intelligence.*

Tiziana Thank you Barbara for this opportunity. It's really nice to have a chance of discussing AI in an interview at AISoLA. My name is Tiziana Margaria. I am a Professor of Software Systems at the University of Limerick in Ireland. I am also involved in several national research centers on software (Lero), on smart advanced manufacturing (Confirm), and in particular the Centre of Research Training in AI (`https://www.crt-ai.ie`), where I am the director in Limerick. My "trade" is formal methods, so my look at AI is from the perspective of the design and verification of the correctness of systems. I have been using reasoning methods and building reasoning tools for over 30 years, so I am interested in AI developments that align directly with my work. It is really exciting to see what is happening now.

Barbara *Can you name one or two AI-related research questions you are currently working on?*

Tiziana Not so directly: it is more the attempt to find a good place for AI and ML in this field. Some people distinguish between AI and ML, others still tend to conflate the two, in the context of software ecosystems as well as of socio-technical ecosystems. Right now, in education and research, I am concerned with the bigger picture. AI is one of the means that help humans to reason, predict or classify. They cover a specific role within a much bigger ecosystem: they are components embedded in a larger workflow or end-to-end process that we call the Digital Thread. The Digital Thread is the integration layer of large heterogeneous systems: it weaves data, sensing, decisions, and actuation, combining AI/ML software with other software, with cyber-physical systems such as machinery, sensors and Internet of Things, telecommunication networks, and so forth, to solve the problem at hand [16]. The thread weaves systems that produce data with the components that analyse the data, where AI has a role, producing outputs that are useful to steer the system, or to inform decisions, or to display information, as in many applications in businesses and society [10, 5]. The role and challenges I see are essentially guaranteeing interoperability [13]: all these tools, including AI/ML ones and their models, must be able to interoperate with other tools, considering dependencies, compatibility, as well as the ability to exchange them with others, guaranteeing equivalence or guaranteeing certain differences. The nature of the application also makes a difference: high assurance software and systems come with financial responsibility or responsibility for the lives of people, for example in the transportation or medical domains. Other applications, like entertainment or advertisement, do not have the same degree of responsibility and regulation. Considering recommender systems, different application profiles require different levels of confidence that the recommendation suits the case at hand. Whether a traditional controller or an AI component,

the recommendations for high assurance contexts must be trustable and justifiable, so that explainability becomes a crucial property. I am working on different concepts of explainability, that depend on the kind of application and context in which it operates: explaining choices [11], classifications/decisions [9], or outcomes of synthesis [17].

Barbara *Does explainability also depend on the type of user group?*

Tiziana Yes, for example the level of confidence that the technologist has in some kind of outputs is determined by the fact that this technologist is able to interpret the outputs. Outputs in my field are typically not data values, like a temperature or a price, but complex outputs like properties, source code or models. There needs to be significant expertise to make sense of those outputs and competently judge their correctness and adequacy. Considering higher level properties, such systems are typically desired to be deterministic systems. This means that if the system runs several times with the same inputs and context, it produces the same outputs. It is predictable. That is not the case for AI systems: they have variability. Probabilistic systems have been in use for a long time: such systems provide an output, typically a recommendation, a classification or a prediction, with an accuracy higher than 50%, and hopefully much higher in cases where actually one expects determinism. Other systems are inherently chaotic: even minor changes to the inputs can result in drastically different outputs. The key issue is that AI and ML technologies fall into the last two categories: they are perfectly suitable for problems where approximations are adequate (one "guesses" anyway, like in advertising), but not when determinism is sought or required by regulators. They may not be adequate or provide sufficient control when lives are at stake, yet enthusiastic adopters are (knowingly or not) using them nevertheless. Explainability, intended as a justification case by case of what led to a certain output, is here the key and the challenge. These three categories of systems have different ways of approaching an interpretation of outputs, and require competence in understanding when the output is good enough. This distinction is hard to communicate to a large part of the potential users, because this way of thinking is unfamiliar to them, it is not part of their categories of judgement.

"We risk to collectively regress from an age of reason and responsibility, where knowledge and skills are assets, education is expanding and reaching ideally everyone, to a new age of belief and delegation to "black box" authorities: entities that we are unable to question, as we either never learned how they work, or have given up that knowledge, out of "convenience"."

Barbara *What is the role of trust in the adoption of AI?*

Tiziana Excellent question. Trust was a significant topic of discussion this week, but with a very different meaning. My background is in high assurance computing. Traditionally, in this community trust is established through certifiability, responsibility, and even forensics, which concerns providing explanations a pos-

teriori, after something has already happened [8]. For this kind of trust one needs to be able to follow up case by case what led to that recommendation, action, or chain of decisions and actions, that resulted in a catastrophe, incurred high costs, or caused other problems, as in Clayton Christensen's opening scenario in his Innovator's Dilemma book [6]. We are back to the quest for determinism, and the current AI is not at that level. I could conclude that AI is not usable and end the conversation here, but that is not my approach. The question is how we can recover trust in an AI context by combining different reasoning systems, different certification systems, as evidence-providing systems that may provide different elements: one delivers a model, another delivers an explanation, another defines the concept of "adequacy" maybe in terms of degrees of "acceptable" approximation under the given circumstances. For me, it is crucial that explanations are useful to enhance the system. Typically, we seek explanations when things go wrong, like the use of counterexamples during verification and testing. If a test fails under identified conditions, how can we prevent the system from encountering those conditions again? If there are no explanations of just the probabilistic ones like [20] in today's AI and neuromorphic computing, one cannot identify why things went wrong. Such "explanations" are not sufficient. I prefer a neurosymbolic approach, because I think that we need to bring together the capabilities of what we have already known for a long time, in the communities that develop reasoning, dependability, reliability (the symbolic aspects), and the newer take on cognitive systems with AI, ML and bio- or neuromorphic approaches. If we succeed to bring them together in a good way, the neurosymbolic combination can actually make both sides stronger.

Barbara *Do you have any key measures in mind to help ensure ethical AI adoption?*

Tiziana Certainly. From an ethical point of view, AI adoption is often viewed as a process of data consumption. I think that data is just one ingredient, but let's consider data first. One measure of ethical AI is how well one handles imbalanced data sets, for example through the adoption of adequate processes to manage the imbalance [3]. For instance, in the medical field, even prevalent conditions like diabetes affects 7% to 10% of the population, so diabetes detection datasets are still very imbalanced. Training a technique with a data set where the category of concern is very infrequent can lead to problems due to insufficient representatives. Developing good technologies to deal with imbalance is the more crucial with increasing degrees of rarity. The general public has developed concerns about human recognition techniques and potential biases in AI systems that classify or decide. The worry that certain segments of the population or customer base may be disadvantaged due to these biases in the dataset can be addressed by rebalancing the data set and developing robust, adequate approaches to recognize and handle these rare events. On the other hand, there is also the potential for unethical use of otherwise correct and fair AI, which is a bigger problem. Once the methods exist, and they are integrated into full workflows, any technology can be used for good or bad. Unfortunately, what is

considered good or bad largely depends on the societal norms, or professional norms. Some communities are more prone to objection than others, for example in terms of privacy, explicit consent, and exceptions to them [2]. The new European act that is going to be introduced soon will provide guidelines for the producers of AI systems and applications that use AI. I welcome that it distinguishes several levels of concern and levels of evidence required to certify or legally accept these systems for practical use. It shows understanding that the same algorithm may be acceptable for a recommender system but it may not be suitable for recommending a treatment in the case of cancer, where the consequences of imprecision or errors could be severe.

Barbara *So, in terms of the future technical capabilities of AI, on a scale of 1 to 10, where 1 stands for today's artificial intelligence systems like ChatGPT, and 10 stands for artificial general intelligence that surpasses human capabilities. What do you think is possible?*

Tiziana That's a very difficult question. Not long ago, we would not have considered ChatGPT to be a 1 but even beyond this entire scale. Now that it suddenly exists, it is re-scaling everything. I hope that within the realm of AI and machine learning we will keep developing systems that are well-understood, localized, and controllable, and not just race towards the kind of general AI that many worry about. I hope that we will stabilize around maybe five to seven. Not because generalized AI is inherently bad, but because it could easily transform into or be used as something harmful to people, the environment, or any other aspect of the context. If humans lose control over these systems, and these systems become too independent and powerful, we could be in trouble. I hope for a five to seven in the sense that I would like to have really efficient, reliable and explainable systems that augment human capabilities. If this way we can become stronger, faster, see better, see earlier, this could provide expert advice to doctors, field workers, scientists, or artists. That would be fantastic. I would accept to go much further with augmentation than with substitution. Substitution has been discussed in contexts where AI and AI-steered robots may replace workers. While some of those tasks and jobs are onerous or dangerous, this could lead to societal issues because certain job categories would no longer be filled by humans, opening the question of what these humans would then do instead. The substitution of knowledge would be much more problematic. Self-learning and self-evolving algorithms that operate unsupervised by humans could pose great risks unless there are very clear legal guidelines that are inherent in the system and cannot be overridden. However, that's difficult to decide and challenging to implement. There's always the possibility that something just goes astray. Actually, I am near certain that unexpected big problems would arise, like second-order unintended consequences down the line, and I doubt that we would be able to "stop the system" and correct it. That's the part I would prefer society not to experience.

Barbara *There are a lot of different AI-driven futures envisioned these days, from dystopia to utopia. Where would you place yourself?*

Tiziana I am an engineer, that's my DNA. My way of looking at the world, at concepts and ideas is that they are parts of systems, systems of systems, purposes. I think in terms of inputs, outputs and context. In terms of trajectories to the future, if we live in a world where certain resources are becoming scarce, we might be able to solve the energy problem by developing and adopting alternative technologies. Seeing all the news of what is being tried, leading to discoveries and to processes of production that scale, I am quite optimistic about addressing this challenge: we are discovering new inputs and produce new outputs or more of certain outputs in new ways. On the pollution challenge I am also quite positive because we start to view globally excess waste as potential raw materials, and we start to act upon it. I do not yet see a close solution for issues like agriculture and water scarcity. If AI systems can help us use these scarce resources more efficiently, more fairly, and distribute them to a larger segment of the population across different societies, equalizing core imbalances across the north and the south of the globe, then I would welcome them. In this case we might co-develop a utopian society positively supported by AI and by its capability of micro-regulating, quick to take and enforce decisions for the local and global good. But if things go wrong, and these capabilities fall into the wrong hands, be they corporations, governments, or other entities, we could easily fall into dystopia. The skyrocketing price of energy, or water, or food could be manipulated, even when its costs go down. Or there could be supply chains diversion for scarce goods, or it could introduce a level of surveillance that impacts freedom and privacy. Misuse would lead to increased imbalances, to a growing divide between the haves and the have-nots of resources, knowledge, or power. This could potentially lead to conflict, even armed conflicts, resulting in losses for everyone. We have not yet sufficiently internalized the reality that the planet is a closed system, where no deep change can compartmentalize and ringfence its effects. Everything is more or less connected, and we can only manage to tailor, or fine-tune, the intensity of some of these dependencies. My wish is that we can achieve the utopian solution and that we, as a global society, are going to be able to agree not to knowingly run or unknowingly slide into a dystopian situation in the future. The danger of "substitution" of control is the danger that initially ok situations may slide into dystopia, and we do not notice until it is too late, and insurmountable premises for conflict have been reached. So how can we find rules that we can all agree upon, and make sure that they are going to be implemented globally? In my system engineering terms, we now live in a global society, i.e., with global repercussions of even local decisions. It is a closed system, where anything has dependencies and therefore any change (including the passing of time) has an effect and perturbates the status quo. Companies operate globally, our economy is interconnected, and our societies are linked via communication and also migration. As we have seen recently in many disruptive examples, what happens in one part of the world can have significant consequences elsewhere. Therefore, it is pointless to have regulations that only apply to certain regions, like the EU. I believe we need a planetary agreement.

"The key is going to be the regulations, with the lawyers talking to the other disciplines because the will of the society needs to be reflected in the laws, regulations, in the standards and norms."

Barbara *Reflecting on the past few days, we've had a diverse range of interdisciplinary discussions and sessions. Was there a particular insight that stood out to you?*

Tiziana My key insight is the fact that we attribute a different meaning to the same words. For example, the degree of precision that is attached to certain words like trust, privacy, and more in a technical context is not shared by attendants from other contexts. We had digital humanists with a discursive approach to subject matters, and then computer scientists and formal methods people who are on the mathematical and rigorous side of system analysis and design. When a concept like trust comes up in the context of ethics, or rules and regulation, or governance, we mean different things by the same word. When we think about governance in computer science, we mostly do not think of asking the population or surveying opinions and preferences. For us, governance is a set of hierarchical decision systems, where there are some rules and there is a technical system that conforms to these rules and enforces them case by case. It is not easy to talk to people who have a less structured and defined understanding of the same words in a way that makes sense to us. So we could communicate more, and in order for it to be meaningful and useful, we should agree on the meaning of certain vocabulary, that becomes a joint and common vocabulary. In other words, there is still fundamental work ahead before we can start working together with less frustration. The second insight is that hybrid tracks like AI in medicine and the digital humanities track did meaningfully involve medicine, history, geography and other disciplines. Here I liked the demonstration of the fact that computer scientists and experts of other disciplines indeed succeeded to find a meeting ground. So, yes we can! There, computer science achieved its best: being a science with its own depth and challenges, which is inherently an application domain-agnostic art. In my opinion computer science is today the most fundamental discipline, as it influences and supports the progress of every domain. I also believe that its impact is most evident when it helps others succeed. A number of applications were really pleasing because two communities that normally would be distinct were able to collaborate and create something beneficial for both.

"What happens in one part of the world can have significant consequences elsewhere. Therefore, it's pointless to have regulations that only apply to certain regions, like the EU. I believe we need a planetary agreement."

Barbara *Is there a particular research question you'd like to see addressed in an interdisciplinary manner in the future?*

Tiziana There is indeed, but it is a meta-level question rather than a concrete one. It is related to the interdisciplinary ability to carry out a meaningful discourse: it revolves around establishing what is considered "known" or "understood" enough, or "concrete" enough, or "good" enough, beyond the community of computer science and engineering. We have sets of measurable criteria for evidence, like certain kinds of tests and a certain level of test coverage, defined

kinds of approximation and acceptability criteria, ways of conducting formal proofs to establish certainty, or code reviews carried out according to certain standards: this establishes fact-based evidence accepted in our community. I am not aware of similar criteria in other communities. In medicine, there are procedural standards for conducting patient-related clinical trials, but I am not aware that philosophy or psychology or sociology or law have similar kinds of criteria. The criteria may well be different from discipline to discipline, but it would be important that we establish some, even if coarse, and that we are able to communicate them and talk in those terms across communities. Otherwise we risk to never be able to establish a communicable ground, as we experienced this week with the discussion on "trust", because one group is happy with an undefined, abstract or even ambiguous meaning, that those who need precision in order to measure, quantify, and decide, cannot work with, because it is too abstract or too ambiguous a concept. Defining some terms and some criteria for achieving those quality levels would enable statements like "I have understood this", "this is known", "this is well accepted", "this is trustable to a level of X", to be shareable and mean the same across communities.

Barbara *What are the different levels of trust that you think are important? Going back to the topic of trust in general, which has been discussed quite frequently in the last week. Is it about establishing a common language and a common concept of trust, or is it about revealing the different concepts of trust at a sufficiently precise level to be able to translate between the different languages?*

Tiziana This sends me back to 35 years ago, when many in the formal methods community dreamed of a common language for modeling systems. It was a mirage: any modelling language was never good enough, because different subsystems had different needs, the concepts could not be faithfully mapped into the other language, and models would be either redundant (too much unneeded information) or not useful (too abstract, missing information). That's when we founded the STTT journal (see the Editorial [7]) and the ETI platform [19], followed by instances for the FMICS Working Group [4, 15] and for the Bioinformatics or climate change domains [12, 1]. I am a strong supporter of the "archipelago" approach: each island is a community with their own vocabulary and criteria, and they are joined by carefully chosen and well-designed bridges of translation. So a concept in one community may be called differently in the other community, but there are clear criteria for expressing how close they are, in which sense and how much they differ, so that we can effectively map the understanding in a meaningful and usable way.

"I hope [for] the utopian version. I am however very scared that it may go otherwise: one or two rogue instances could be enough to create enormous damage."

Barbara *Do you have a suggestion on how to move towards the archipelagic approach? How can we establish the necessary level of precision and transparency to bridge the gap between disciplines?*

Tiziana I think it is possible in application domains that work with a quantitative approach, though it requires a lot of work. Understanding and mapping the concepts in more discursive fields might be harder as these communities might not be accustomed to the level of concreteness that we need even to start expressing models. In CS as in STEM we are always comparing things, so if there is no definition of what we are comparing, or at least a qualitative measure like high, medium, or low that are rooted in properties of the system and agreed by both communities, it is difficult to have meaningful conversations. I observed this week that some communities are closer to achieving this level of definition than others, and that some do not realize that there is a problem. So we might even have to take one step back, and have to communicate first our need for establishing something of that kind of precision in order for us to become operational. We are always told that we need to have SMART goals, but where is the SMART definition in terms of measurable, achievable, and so on? It remains too vague.

Barbara *What do you think the vision of AI should be?*

Tiziana I see AI as just another set of technologies that have been developed. Before the "dotcom" bubble burst, it was internet everywhere because this was the big novelty at that time, the new technology that would revolutionize everything. It led to the bubble, with a lot of pain and repercussions, but eventually we found ways of using it proficiently, and a way of actually bringing it to everybody with, importantly, more advantages than disadvantages, so that is an innovation wave that has been meanwhile internalized in our society. In previous times, people were scared by transportation with cars, or could not fathom having a telephone in their homes, as for urgent communication there were messenger boys at hand. Now we have essentially a supercomputer in each phone, instant communication is perceived as a basic need and a right, and horse driven carriages are nowadays largely for tourists. Past some hype, I believe that AI has the potential to become the new "household tool" that aids people. Indeed I hope that this is the vision that we are going to enact: the utopian version. I am however very scared that it may go otherwise: one or two rogue AI instances could be enough to create enormous damage. So on one side I am hopeful, because the research and science communities are willing to collaborate, and there is a will to design a shared, positive, and useful AI vision. The problem has always been controlling the outliers: the recent Oscar to the film on Oppenheimer is a timely reminder of what it means to "master" a new technology in the context of entities and decision makers with conflicting purposes and goals.

Barbara *Are there other perspectives you would like to see addressed or integrated into the interdisciplinary discourse? This week we've met with psychologists, lawyers, computer scientists, and philosophers. We've also seen various AI applications in health care and history, as you mentioned. Do you see other disciplines that should get more involved?*

Tiziana The key is going to be the regulations, with the lawyers talking to the other disciplines because the will of the society needs to be reflected in

the different flavours of the legal system: laws, regulations as well as in the standards and norms. The values define the essence of a society: it is from values that we derive consequences like the kind of governance for yourself and others. So I would say that lawyers, engineers, social engineers should be all involved because this is where the value-defining discourse that is already happening with end users, social stakeholders etc., from the philosophical level becomes concretized in something that can be then translated for and into the technical systems. If there is a lot of discourse but it does not find its way into the systems, there is an insurmountable gap between the talk and the walk. My concern goes back again to the digital thread: we need to establish this useful communication thread between the different disciplines, and I think that the regulations and law perspective is the common point where each one of us has the ability to contribute meaningfully and to take out something useful.

Barbara *Is there anything else you would like to add?*

Tiziana I am glad that we had those conversations for this entire week. I am a systematic initiator of weird interdisciplinary tracks at ISoLA and on initiatives where I espouse fundamental values and concerns, see for example the entire EU Strategic Support Action on Simplicity over 10 years ago [14] and my insistence of the relevance of extreme model driven and low-code/no-code approaches for sustainability and democratization for years [18] and even more now in the R@ISE project (`https://software-engineering.ie/raise/`). Accordingly, I find such projects and discussions extremely enriching as a person, as a citizen, as a curious individual and also as a researcher and a producer of IT. I hope that there will be more of this in the future and thank you very much for this opportunity.

Barbara *Thank you, Tiziana, for your time and insights. Your interdisciplinary and application-driven perspectives have been invaluable. Have a great day!*

Tiziana Thank you, Barbara.

References

1. Al-Areqi, S., Lamprecht, AL., Margaria, T. (2016). Constraints-Driven Automatic Geospatial Service Composition: Workflows for the Analysis of Sea-Level Rise Impacts. In: Gervasi, O., et al. Computational Science and Its Applications - ICCSA 2016. Lecture Notes in Computer Science, vol 9788. Springer. `https://doi.org/10.1007/978-3-319-42111-7\_12`
2. B. Alamri, I. T. Javed and T. Margaria, "A GDPR-Compliant Framework for IoT-Based Personal Health Records Using Blockchain," 2021 11th IFIP International Conference on New Technologies, Mobility and Security (NTMS), Paris, France, 2021, pp. 1-5, doi: 10.1109/NTMS49979.2021.9432661.
3. Amandeep Singh, Olga Minguett, Tiziana Margaria, Binary Decision Diagrams and Composite Classifiers for Analysis of Imbalanced Medical Datasets, ECEASST Volume 82: 11th International Symposium on Leveraging Applications of Formal

Methods, Verification and Validation - Doctoral Symposium, 2022, ISSN 1863-2122, http://dx.doi.org/10.14279/tuj.eceasst.82.1227.1147
4. Arenas, Alvaro E., Bicarregui, Juan C., and Margaria, Tiziana. 'The FMICS View On The Verified Software Repository'. Journal of Integrated Design & Process Science, vol. 10, no. 4, pp. 47-54, 2006
5. Chaudhary, H.A.A. et al. (2022). Model-Driven Engineering in Digital Thread Platforms: A Practical Use Case and Future Challenges. In: Margaria, T., Steffen, B. (eds) Leveraging Applications of Formal Methods, Verification and Validation. Practice. ISoLA 2022. Lecture Notes in Computer Science, vol 13704. Springer. https://doi.org/10.1007/978-3-031-19762-8_14
6. Christensen, Clayton M. 2016. The Innovator's Dilemma. Management of Innovation and Change. Boston, MA: Harvard Business Review Press.
7. W. Rance Cleaveland, Tiziana Margaria, Bernhard Steffen, Editorial, International Journal on Software Tools for Technology Transfer (STTT) Volume 1 Issue 1-2 December 1997 pp 1–5, Springer, https://doi.org/10.1007/s100090050001
8. Stefania Gnesi, Tiziana Margaria, Formal Methods for Industrial Critical Systems: A Survey of Applications Wiley-IEEE Computer Society, ISBN: 978-1-118-45987-4, November 2012 Wiley-IEEE Computer Society
9. F. Gossen, T. Margaria and B. Steffen, "Towards Explainability in Machine Learning: The Formal Methods Way," in IT Professional, vol. 22, no. 4, pp. 8-12, 1 July-Aug. 2020, doi: 10.1109/MITP.2020.3005640.
10. Guevara, I.; Ryan, S.; Singh, A.; Brandon, C.; Margaria, T. Edge IoT Prototyping Using Model-Driven Representations: A Use Case for Smart Agriculture. Sensors 2024, 24, 495. https://doi.org/10.3390/s24020495
11. Kubczak, C., Margaria, T., Steffen, B., Winkler, C., Hungar, H. (2009). An Approach to Discovery with miAamics and jABC. In: Petrie, C., Margaria, T., Lausen, H., Zaremba, M. (eds) Semantic Web Services Challenge. Semantic Web And Beyond, vol 8. Springer. https://doi.org/10.1007/978-0-387-72496-6_13
12. Lamprecht, AL., Margaria, T. & Steffen, B. Bio-jETI: a framework for semantics-based service composition. BMC Bioinformatics 10 (Suppl 10), S8 (2009). https://doi.org/10.1186/1471-2105-10-S10-S8
13. Margaria, T., Chaudhary, H.A.A., Guevara, I., Ryan, S., Schieweck, A. (2021). The Interoperability Challenge: Building a Model-Driven Digital Thread Platform for CPS. In: Margaria, T., Steffen, B. (eds) Leveraging Applications of Formal Methods, Verification and Validation. ISoLA 2021. Lecture Notes in Computer Science vol 13036. Springer. https://doi.org/10.1007/978-3-030-89159-6_25
14. T. Margaria, B. D. Floyd and B. Steffen, "IT Simply Works: Simplicity and Embedded Systems Design," 2011 IEEE 35th Annual Computer Software and Applications Conference Workshops, Munich, Germany, 2011, pp. 194-199, doi: 10.1109/COMPSACW.2011.42.
15. T. Margaria, H. Raffelt, B. Steffen and M. Leucker, "The LearnLib in FMICS-jETI," 12th IEEE International Conference on Engineering Complex Computer Systems (ICECCS 2007), Auckland, New Zealand, 2007, pp. 340-352, doi: 10.1109/ICECCS.2007.43.
16. Margaria, T., Schieweck, A. (2019). The Digital Thread in Industry 4.0. In: Ahrendt, W., Tapia Tarifa, S. (eds) Integrated Formal Methods. IFM 2019. Lecture Notes in Computer Science(), vol 11918. Springer. https://doi.org/10.1007/978-3-030-34968-4_1
17. T. Margaria and B. Steffen, "LTL Guided Planning: Revisiting Automatic Tool Composition in ETI," 31st IEEE Software Engineering Workshop (SEW 2007), Columbia, MD, USA, 2007, pp. 214-226, doi: 10.1109/SEW.2007.63.

18. Margaria, T., Steffen, B. (2020). eXtreme Model-Driven Development (XMDD) Technologies as a Hands-On Approach to Software Development Without Coding. In: Tatnall, A. (eds) Encyclopedia of Education and Information Technologies. Springer. https://doi.org/10.1007/978-3-030-10576-1_208
19. Steffen, B., Margaria, T. and Braun, V., 1997. The Electronic Tool Integration platform: concepts and design. International Journal on Software Tools for Technology Transfer, 1(1), pp.9-30.
20. SHAP: (SHapley Additive exPlanations) https://shap.readthedocs.io/en/latest/