NEWSLETTER °2
NOVEMBER 2022

Welcome to the FOCETA project newsletter n°2!

In a context where applications are increasingly being developed based on complex autonomous systems driven by artificial intelligence, their safety, autonomy and trustworthiness are challenging, especially those of learning-enabled systems, not easily traced by continuous engineering.

The ultimate goal of the FOCETA project is to develop the foundations for continuous engineering of trustworthy learning-enabled autonomous systems. The underlying targeted scientific breakthrough of FOCETA lies in the convergence of model-driven and data-driven approaches. This convergence is further complicated by the need to apply verification and validation incrementally and avoid complete re-verification and re-validation efforts.

Our public newsletters will keep you up-to-date on the progress made within the project. You will discover how the consortium partners cooperate to achieve the project objectives. You will also know how and when we disseminate the FOCETA results.

THE FOCETA PROJECT IN A NUTSHELL

- **4 985 540 € EU funding**
- **656 Person-months**
- **An international consortium**
- **36 months project duration**
- **Started on 01/10/2020**
- **13 partners from 8 countries**

Read more about FOCETA on the project website

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

EDITORIAL BY THE PROJECT COORDINATOR

Using Artificial Intelligence (AI) in the form of machine learning has the potential to modify many safety-critical applications, such as those in automotive and healthcare. Nevertheless, despite notable investment and excellent demonstrations, such technologies have struggled to live up to their pledges. In order to better focus our efforts, to find the right strategy for this investment, we must first ask the right questions to define the right research directions. Of course, these questions can vary and depend on the research topic. Nevertheless, the three questions that seem essential to someone interested in rigorous systems design are:

1. Is AI in mission-critical systems hype, or is it something sustainable?

The application of artificial intelligence is a reality in many fields, including critical applications such as transport or medicine. The maturity of these technologies has reached a point where they are now clearly offering benefits through various safety applications. Unfortunately, the use of AI components that do not provide a guarantee of trust can lead to dramatic situations. For example, the Tesla crash and the death of a pedestrian in the Uber car accident were caused by computer vision. Of course, one can adopt a conservative approach that can be easily justified, recommending the non-use of AI components in critical applications. However, this conservative measure cannot be sustained over time as more and more manufacturers are pushing to take advantage of the benefits of AI in safety applications. Therefore, the forbidding of integrating AI components in critical systems cannot stand up to the innovative applications of these technologies. Hence, engineering and scientific foundations are needed for designing critical applications integrating AI components.

2. What is the difference between the engineering of pure software and learning-enabled systems, especially mission-critical ones, such as autonomous learning enabled (LE) systems?

Most of the significant differences in development methodology between AI systems and software-based systems are due to the following:

- The central role played by data in AI systems. The collection, cleaning, management, and continuous data update add new tasks.
- Uncertainty is the dominant characteristic in AI-enabled systems.
- Increased urgency of making progress on how to model, analyze and safeguard against the inherent uncertainty of our systems.
- AI specifications are specifications of problems, not the behavior of systems.
- Verification challenges are inevitably exacerbated in AI-enabled systems, given their inherent uncertainty.
- The continuous update is a big challenge in AI-enabled systems.
- We know the challenges of designing embedded systems that rely on integrating many disparate SW/HW components. AI components developed independently are yet another set of subcomponents which behavior needs to be reliably predicted.

3. How can we ensure the trustworthiness of these critical systems with AI components?

Trust is imperative for AI widespread adoption and, thus, overall success. Establishing trust, however, is not a straightforward process. Several constituent elements of trust cover all the dimensions of an AI system. The criteria catalog we can find in the literature to improve safety in critical applications can be summarized as the following:

- All algorithms based on decision-making must be explainable.
- The functionality of algorithms must be analyzed and validated using formal verification methods before use.
- Statistical validation is necessary, mainly in cases where formal verification is unsuitable for specific application scenarios due to scalability issues.
- The inherent uncertainty of neural network decisions must also be quantified.
- Systems must be observed during operation, for example, by using online monitoring processes.

Prof. Saddek Bensalem
Université Grenoble Alpes/Verimag
France

NEWS & EVENTS FROM FOCETA

The **joint H2020 FOCETA - OpenDR Summer School** tooke place on 3-7 October 2022 in Thessaloniki >> Feedback article in the Get Together section on p. 10

38 scientific publications are available in open access on the project website >> Read

Watch the FOCETA project **video** >> Watch

Download the FOCETA project **leaflet** >> Download

Download the FOCETA project **poster** >> Download

Download the FOCETA **roll-up banner** >> Download

PROJECT CONTACT INFORMATION

Website: http://www.foceta-project.eu/

Project Coordinator: Saddek Bensalem (Université Grenoble Alpes)

Technical Coordinator: Dejan Nickovic (AITAustrian Institute of Technology)

PMO/Dissemination: Sofia Santi (L-UP SAS)

FOCETA

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

NEWS FROM FOCETA

THE FOCETA DESIGN FLOW METHODOLOGY

With the tremendous advances in artificial intelligence (AI) and machine learning (ML), learning-enabled autonomous systems (LEAS) are revolutionizing today's society with applications ranging from self-driving vehicles to smart medical devices. To enable their wide acceptance by the users and the regulatory bodies, LEAS must conform to strict correctness and safety requirements. This is nevertheless very challenging to achieve due to the interplay between learning-enabled (LE) and classically engineered components, operating together in sophisticated and unpredictable environments. Despite many recent efforts to address various verification and validation aspects of safety in autonomy, there is a lack of systematic design methods for LEAS that ensure their correctness and safety.

We develop in FOCETA a holistic methodology for designing correct and safe LEAS through bridging the gap between the currently applied development, verification, and validation techniques for LE systems and their operation in the real world. The proposed engineering approach also considers the foreseen updates to LE components, in response to emerging requirements from new scenarios, imperfect knowledge of the machine learning models (noise in data observations) or contextual misbehavior of them and possible security threats. This need is addressed within a continuous development and testing process mixing software development and system operations in iterative cycles. In this process, formal specifications play an important role and are used throughout the system's life-cycle. They are used to remove ambiguities in the exchange of knowledge and information, to guide formal verification and testing activities, and to monitor the LEAS during operation for requirements violation.

The FOCETA methodology is based on two design flows. The first is to develop trustworthy LE components, and the second is to ensure correct and safe LEAS, through iterative cycles of development and system operations. Our methodology is based on the verification of simulation results and the collection of additional data for system improvement. Runtime monitoring techniques are used to supervise the decision-making of the systems and record environmental situations.

The rigorous design approach used to develop a safe and trustworthy LE component heavily depends on the component's intended functionality. In FOCETA, we consider mainly two classes of LE components - perception modules and data-driven controllers. Designing a trustworthy perception module is very different from designing a safe data-driven controller. While an object detection component based on a neural network (NN) does not admit a natural mathematical characterization of its expected behavior, the correctness and safety requirements of a data-driven controller can be described using formal specifications. For the perception components, we refine the usual development cycle consisting of training, verification and validation, and deployment with additional activities based on runtime monitoring, explainable AI and reliability estimation to support the certification of the component. For the data driven controllers, FOCETA explores specification-driven safe reinforcement learning and runtime enforcement techniques.

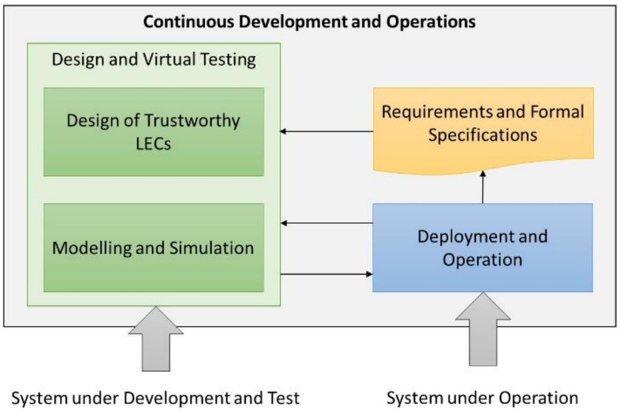

Figure 1: FOCETA Methodology flow.

The design flow for LEAS consists of two parts: the design part (design time) for the design, implementation, and verification of LEAS, and the (run time) operations part, focusing on the deployment of the LEAS and their operation in the real world. The current state of practice is extended towards the transfer of knowledge about systems and their contexts (e.g., traffic situations) from the development to operations and from the operations back to development, in iterative steps of continuous improvements. Over the complete life cycle of autonomous LE systems - from specification, design, implementation, and verification to the operation in the real world – the methodology enables their continuous engineering with a particular focus on the correctness with respect to an evolving set of requirements and the systems' safety. Moreover, the whole design flow ensures traceability between requirements and the system/component design. A key feature is the usage of runtime monitors for the seamless integration of development and operations. Runtime monitors observe a system (part) via defined interfaces and evaluate predefined conditions and invariants about the system behavior based on data from these interfaces. This allows identifying the need for LE component updates during continuous testing/verification, if, for example, some data in a test scenario result in a safety property violation or if a new requirement will emerge in response to a previously unknown adversarial threat.

Written by Dejan Nickovic (AIT Austrian Institute of Technology)

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

KEY RESULTS FROM THE FOCETA CORE SCIENTIFIC WORK PACKAGES AND HOW THEY FIT INTO THE FOCETA DESIGN FLOW METHODOLOGY

As explained in the FOCETA Newsletter n°1 (article "The FOCETA project methodology"), the FOCETA project relies on three key scientific work packages, respectively (i) "Modelling and simulation for autonomy", (ii) "Verification and validation for autonomy", and (iii) "Agents for Performance and Security beyond safety via Runtime Monitoring and Enforcement". These work packages originate most of the methods and tools produced in the project that will be applied to the two FOCETA industrial use cases, i.e., automotive valet parking and smart medical devices.

In this article, we illustrate one prominent key result of each of these three work packages and explain how they fit into the FOCETA design flow methodology.

KEY RESULT FROM THE WORK PACKAGE "MODELLING AND SIMULATION FOR AUTONOMY": A DIGITAL TWIN ARCHITECTURE WITH FORMAL ANALYSIS CAPABILITIES FOR LEARNING-ENABLED AUTONOMOUS SYSTEMS

Learning-enabled autonomous systems, such as self-driving cars, combine sensor components with automated controllers for advanced decision making, include learning-enabled components (LECs) and operate in complex and unpredictable environments. Their safety will have to be assessed with respect to an ever-evolving set of (critical) scenarios, instantiating difficult conditions (e.g., traffic conditions) for their sensors and algorithms. This scenario-based approach assumes a progressive and iterative development of the learning-enabled system, for gradually exploring the operational conditions, in which the system can safely operate (Operational Design Domain - ODD). In this process, we have to take into account a set of continuously evolving requirements and the need for continuous testing/verification. Moreover, we also have to support possible LEC updates in response to emerging requirements from new scenarios, imperfect knowledge of the machine learning models (noise in data observations) or contextual misbehavior of them and possible security threats.

These challenges are addressed in FOCETA work package "Modelling and simulation for Autonomy" through continuous development and testing by mixing software development and system operations in iterative cycles, which alternate between simulation-based verification/validation and physical system testing (operations). Simulation-based analysis is a cost-effective means to verify the system's performance over diverse parameter ranges and massive sets of scenarios, including scenarios that may be impossible to realize when the system operates in its environment (e.g., with very rare events).

FOCETA partners involved in the work package "Modelling and simulation for Autonomy" have invested on a multi-modeling and compositional simulation-based digital twin architecture. Multi-modeling means that it is possible to mix model components at various abstraction levels that have been designed with different tools and frameworks (e.g., discrete modeling for the system's computing elements and continuous modeling for its physical components). Our architecture (Figure 2) allows building a virtual model, for analyzing and diagnosing the system's operation in simulated time (digital twin) and supports model-in-the-loop (MIL), software-in-the-loop (SIL) and hardware-in-the-loop (HIL) verification/validation of the system's functions, at various abstraction levels.

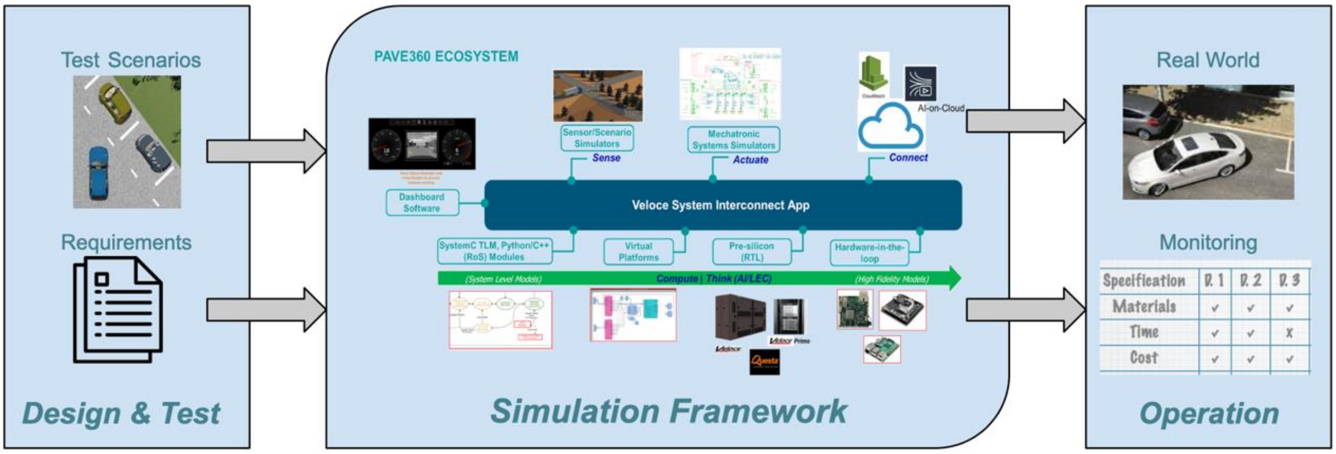

Specification	R 1	R 2	R 3
Materials	✓	✓	✓
Time	✓	✓	x
Cost	✓	✓	✓

Figure 2 : Compositional simulation-based digital twin architecture for virtual testing of learning-enabled autonomous systems.

The digital twin architecture is based on the PAVE360-VSI interconnect fabric, provided by FOCETA partner Siemens Industry Software, that was extended through integrating formal modeling capabilities using the BIP (Behavior, Interaction, Priority) component framework developed by Université Grenoble Alpes. The BIP extension allows mixing model-based and machine learning components, in order to support the design, training, verification and potential update of LECs.

The PAVE360-VSI interconnect fabric features cyber-physical ports that enable heterogeneous client connections, including sensor/scenario simulators (e.g., Simcenter Amesim by Siemens), mechatronic system simulators (e.g., Simcenter Prescan by Siemens), co-simulated functional mockup units, cloud services and computational models (e.g., in C/C++, SystemC TLM, Python, Robot Operating System), virtual platforms and hardware emulation modules. In essence, PAVE360-VSI acts as the core of our compositional architecture, in which simulation progresses in discrete time steps; it keeps the simulated time and is responsible for all time advance operations that are coordinated across the client processes.

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

Online formal analysis of simulation traces is also supported by Aristotle University of Thessaloniki for verifying safety properties for the overall system model. This takes place through the integration of runtime monitors, generated by the DejaVU monitor synthesis tool, from formal specifications in the past-time fragment of first-order Linear Temporal Logic (LTL). The latter was introduced in the FOCETA project work package "Industrial requirements and incremental safety-and-security cases". The past-time first-order LTL is the target language of the tool-supported FOCETA requirements specification/formalization approach developed by Aristotle University of Thessaloniki. This approach features ontology-based semantic modeling and analysis capabilities, for assessing the quality (consistency, completeness, clarity, etc.) of a set of continuously evolving natural language requirements, such as those seen when mixing software development and scenario-based (physical or virtual) testing in iterative cycles.

These key results aim to provide the gluing context for integrating the technological developments in the other FOCETA work packages which contribute to the design of trustable machine learning components, as well as to the efficient deployment and operation of learning-enabled autonomous systems.

KEY RESULT FROM THE WORK PACKAGE "VERIFICATION AND VALIDATION FOR AUTONOMY": A SET OF TECHNIQUES FOR SCALABLE AND RIGOROUS ANALYSIS OF DEEP LEARNING BASED PERCEPTION AND CONTROL COMPONENTS

One of the key objectives of the FOCETA project work package "Verification and Validation for Autonomy" is to develop practical verification and testing techniques for the perception and control components of autonomous systems that are learned from massive data. Previous work on the verification and testing have been popular but with gaps to be filled. For example, verification techniques have been focused on small models, and testing techniques rely on coverage metrics that lack systematic justifications on their relationship with the properties to be tested. Neither of them considers the context where the machine learning model is running.

The efforts done in this work package over the past year were to address these gaps. For example, consortium partner DENSO has developed a novel robustness verification method for attention networks; consortium partner University of Liverpool has developed a model agnostic reachability analysis method. Both methods can work with neural networks that are significantly more complex and larger than the usual verification techniques can deal with. Moreover, multiple works have been conducted by the FOCETA partners to identify realistic problems, such as Intel's work on bug detection for object detection models and SIEMENS' work on sensor fault-detection and sensor fusion for robust and reliable sensing. In the following, we highlight a joint effort between the University of Liverpool and the FOCETA Project Coordinator, Université Grenoble Alpes.

It has been noticed that for practical applications, adversarial examples, i.e., defects of robustness property, are meaningful only when they cannot be differentiated through human's perception and are on or close to the input distribution. University of Liverpool has proposed a novel hierarchical distribution-aware testing that is able to detect meaningful adversarial examples, as well as a reliability assessment model, to realistically evaluate the safety performance of a machine learning model in operational time. Specifically, we synthesise the operational profile (formalised as a high-dimensional distribution) from not only dataset but also the expert's prior knowledge and prioritise the finding of adversarial examples and the reliability assessment according to the operational profile.

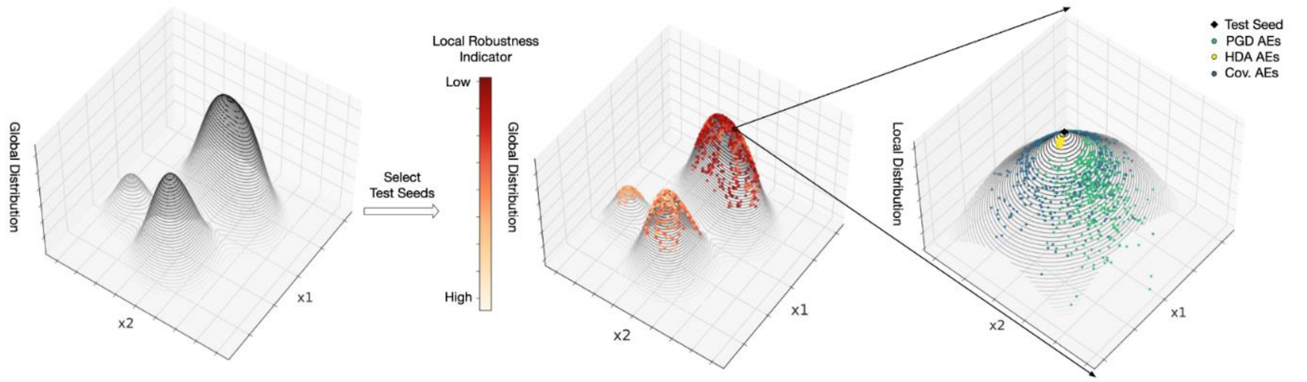

Figure 3: An example of Hierarchical Distribution Aware (HDA) Testing for detecting adversarial examples (AEs), compared with state-of-the-arts including Projected Gradient Descent (PGD) attack and Coverage-guided (Cov.) testing. HDA proceeds in three stages. Stage 1: HDA explicitly approximates the global distribution of input features (projected to a 2D space for illustration); Stage 2: it selects test seeds based on the global distribution and local robustness indicators; Stage 3: it generates test cases around test seeds considering the local distribution (representing pixel-level perception quality) and prediction loss of AEs

KEY RESULT FROM THE WORK PACKAGE "AGENTS FOR "PERFORMANCE AND SECURITY BEYOND SAFETY" VIA RUNTIME MONITORING AND ENFORCEMENT". RUNTIME MONITORING NEURON ACTIVATION PATTERNS AND BEYOND

For using neural networks (NNs) in safety-critical applications, it is crucial to know if a decision made by a NN is supported by prior similarities in training. One of the key results of this work package is to enhance the state-of-the-art in monitoring neuron activation patterns. Figure 4 gives the high-level workfkow on runtime monitoring of on-off neuron activiation patterns. Conceptually, after the standard training process, one creates a monitor by feeding the training data to the network again in

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

order to store the neuron activation patterns in an abstract form. During operation, a prediction over an input is supplemented by examining if the neuron activation pattern is contained in the abstraction. To illustrate, in Figure 4, the scooter is classified as a car, and as its neuron activation pattern is not among the existing patterns created from the training data, the monitor reports that the decision made by the NN can be problematic.

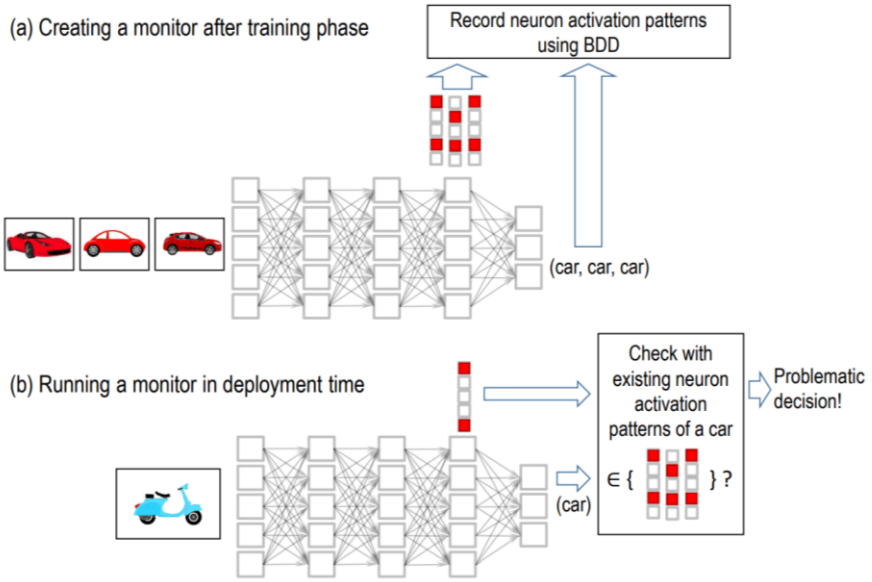

Figure 4: Workflow on runtime monitoring on-off neuron activation patterns.

Within FOCETA, we proposed a novel framework of runtime monitoring of neural networks using box-based abstractions. Box-based abstraction consists in representing a set of values by its minimal and maximal values in each dimension. To build the monitor, we use both the good and bad network behaviours (correct and incorrect decisions respectively) as references to build box abstractions. Since the box-shape abstractions of the good and bad reference behaviour may intersect or not, a new generated pattern can belong to both abstractions.

Our runtime monitor, as illustrated in Figure 5, assigns verdicts to a new input as follows:

- if the input generates patterns that fall only within the good references, then the input is accepted;
- if the input generates patterns captured by both the good and bad references, it marks the input as uncertain;
- otherwise, the input is rejected.

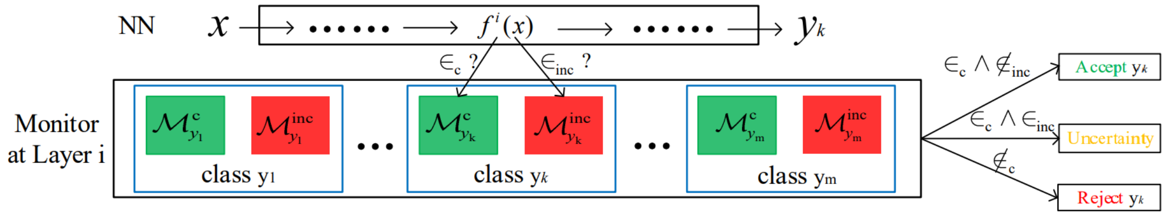

Figure 5: Framework of runtime monitoring of neural networks.

Introducing uncertainty verdicts allows identifying suspicious regions when the abstractions of good and bad references overlap. By reducing the abstraction size, one may remove the overlapping regions and obtain suitable abstraction size. Otherwise, it indicates that the network does not have a good separability: the positive and negative samples are tangled. Both cases are illustrated in Figure 6. This provides feedback to the network designer. It also permits comparing the regions of patterns and thus enables the study of the relationship between good and bad behavior patterns of the network.

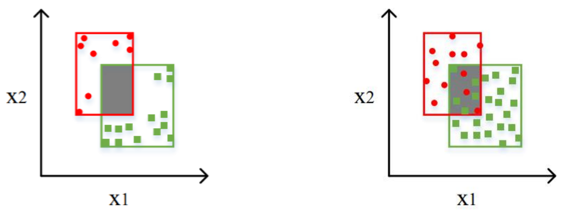

Figure 6: An illustration of the source of uncertainty. (Left) Too coarse abstractions. (Right) Bad separability of classification.

Written by Panagiotis Katsaros (Aristotle University of Thessaloniki), Xiaowei Huang (University of Liverpool) and Bettina Könighofer (Graz University of Technology)

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

THE FLOOR TO THE PLAYERS: INTERVIEW WITH FOCETA PARTNERS

DOUBLE INTERVIEW WITH DR. BETTINA KÖNIGHOFER, ASSISTANT PROFESSOR OF COMPUTER SCIENCE AT THE GRAZ UNIVERSITY OF TECHNOLOGY, AND DR. SON TONG, SENIOR RESEARCH ENGINEER AND PROJECT MANAGER AT SIEMENS DIGITAL INDUSTRIES

***Dr. Bettina Könighofer** holds a PhD degree in Computer Science from the Graz University of Technology. Bettina's research interests lie primarily in the area of reinforcement learning, formal verification, and runtime verification and enforcement. Today, Bettina is assistant professor of computer science at the Institute of Applied Information Processing and Communications of the Graz University of Technology. Within FOCETA, she is leading the work package entitled "Agents for "Performance and Security beyond Safety" via Runtime Monitoring and Enforcement".*

***Dr. Son Tong** obtained his PhD in learning control for mechatronics systems from KULeuven (Belgium) in 2016. Currently, Dr. Tong is a senior researcher and leading a R&D team of research engineers and industrial PhDs at Siemens Digital Industries Software, working on engineering projects in autonomous driving, control, and AI topics. Within FOCETA, Son is the leader of the work package "Methods and tools applied to industrial cases demonstrations".*

Figure 7: Bettina Könighofer (Graz University Of Technology) and Son Tong (Siemens Digital Industries).

Question 1 (Q1): Bettina, how did you become passionate about computer science? And Son, how did you become passionate about control, AI, and autonomous driving?

Answer 1 (A1) / Bettina: I liked computer science from my programming courses at the technical high school, but I got passionate about theoretical computer science during my studies. It is intriguing how many challenging problems exist in the field, many of which are undecidable or of such high complexity that modern computers are only able to solve small problem instances. But still, computer scientists do not get intimidated, and develop intelligent algorithms and high-efficient tools to tackle these challenges. I am passionate about finding new algorithms for problems which are condemned to be practically unsolvable because of their high complexity.

A1 / Son: I am keen on developing and optimizing systems toward higher performance, higher efficiency, environment-friendly and smarter. They range from small things like house appliances to mechatronics systems in manufacturing and transportation systems. Control is a great tool to do that without the need of changing hardware. The mechanism of continuous feedback, feedforward, and prediction via rigorous system analysis and computation is very efficient to optimize systems for specific objectives. Moreover, artificial intelligence (AI) technologies look at another way to improve systems via exploiting data, which I find complementary to conventional control systems. Recently, I got motivated to tackle complexity and multidisciplinary challenges in smart transportation and mobility revolution, with the aim to make driving more safe, comfortable and eco-friendly.

Q2: Bettina, within the FOCETA project, you are the leader of the work package entitled "Agents for 'Performance and Security beyond Safety' via Runtime Monitoring and Enforcement". What are the objectives and challenges of this work package?

A2 / Bettina: The goal of this work package is to design intelligent, high-performance, and transparent learning agents with guaranteed correct behaviour. The greatest challenge that we are facing is that we want to have safety guarantees for systems that are incredibly complex, which have a very high-dimensional input and state space, and the networks lack any structure that we can exploit. This renders a full formal verification of the systems practically unfeasible. For that reason, in our work package we follow the approach to monitor the system execution during runtime, and to detect any unsafe or incorrect behaviour early enough to prevent any harm from happening.

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

Q3: Within this work package, both Graz University of Technology and Siemens Digital Industries are contributing towards the achievement of the currently ongoing task "Runtime enforcement for ensuring performance while guaranteeing safety and security". How would you summarise your work in this task and how do you collaborate with each other?

A3 / Bettina: Graz University of Technology develops new methods to automatically synthesize runtime enforcers, the so-called shields. A shield extends a monitor in a way that it not only detects unsafe behavior but also corrects it during runtime such that safe system execution can be guaranteed. Furthermore, shields are minimally interfering with the system such that we can provide safety guarantees while ensuring high performance. A shield is automatically synthesized from a given model of the system and a formal safety specification. Therefore, we say that shields are correct-by-construction. Within this work package, Graz University of Technology tems is working together with several FOCETA partners, including Siemens Digital Industries, to further develop the concept of shielding. We are advancing current shield synthesis approaches, both from the theory side - by developing new algorithms to constructs shields that are able to handle the requirements we need for the FOCETA use cases, as well as from the practical side - by developing highly efficient shield synthesis tools.

A3 / Son: In this work package, Siemens Digital Industries team focuses on two topics in order to deal with safety and performance of learning control systems. First, how we can exploit performance-oriented reduced vehicle powertrain models to be used online in a modular way for system safety monitoring and prediction. The reduced models are in the form of linearized systems or neural networks with light computation, to be deployed online and to capture a sufficient system dynamics performance. Second, we are developing control technologies that can combine conventional model-based control (i.e., optimal control) and AI data-driven based control (i.e., imitation learning, reinforcement learning). The main motivation is to exploit performance from both approaches, while still guaranteeing safety in the design.

Regarding collaboration between FOCETA partners, we have regular meetings to exchange on technical updates and common applications. To mention one example, Siemens Digital Industries provided licenses to consortium partners to use our software tool in the FOCETA project, that we believe will be helpful to accelerate developments and collaborations from both academic and industry sides.

Q4: Another task in this work package is the one dedicated to "Tool and Early Technology Evaluation", under Graz University of Technology lead and with Siemens Digital Industries contribution. Could you explain to what extent this task is key towards the achievement of the whole work package? How do you contribute to its work?

A4 / Bettina: We have several academic partners working together and proposing many new theoretical approaches and propotype implementations for their new methods. The goal of the early technology evaluation is to give the industrial partners a good understanding of the advantages, the potential and the limitations of the proposed methods and tools. Based on this early evaluation, industrial partners will have a clear vision on how these technologies could be concretely applied in the final demonstrating use cases, i.e., in urban traffic autonomous driving and smart medical devices.

A4 / Son: While there are many tools, methods, and algorithms being developed in the project from multiple work packages and partners, it is essential that they are being constantly reviewed by the project consortium partners. The evaluation criteria are for example the maturity level, state of art, performance and practical aspects like interfacing with other tools - especially for the use case integrated demonstrators. In this task, we at Siemens Digital Industries evaluate the tool developments from our industrial perspective. We try to provide insights on how they are relevant to the FOCETA project use cases, and at the same time feedback that can be helpful to our partners in their tool development.

Q5: Son, within the FOCETA project, you are the leader of the work package entitled "Methods and tools applied to industrial cases demonstrations". What are the objectives and challenges of this work package? Could you explain in what extent this work package plays a key role in view of the FOCETA project outcomes? What are Siemens Digital Industries' key assets to lead these activities?

A5 / Son: In this work package, the FOCETA consortium partners work together on final industrial validation and demonstration of the methods and tools developed in the project. In particular, we focus on high technology readiness level at industrially relevant environment. Hence, the work package plays a critical role in the success of the project.

The objective is to showcase the performance and advantage of the new tools; furthermore, how the tools are integrated and thoroughly tested at system level. For industrial ready level applications, safety and robustness are key aspects; hence we aim at testing the system with different scenarios and environments, including critical ones such as adversarial scenarios. The main challenges are how to integrate all the tools efficiently following standard industry interfaces. This collaboration between partners requires intense communication and a great effort.

Within this work package, a key part of the job is done by the use case leaders, i.e., DENSO AUTOMOTIVE Deutschland GmbH on the Automated Valet Parking use case, and R G B MEDICAL DEVICES SA on the medical demonstration use case. My role is to follow the overall progress, timelines, connect with partners from other work packages, and provide feedback when necessary. Besides, Siemens Digital Industries is actively contrubuting on both use cases, since our team has quite some industrial and commercial solution background which can contribute to the achievement of the task objective.

Q6: In line with the FOCETA project methodology, the outcomes of the work package led by Bettina, i.e. "Agents for "Performance and Security beyond Safety" via Runtime Monitoring and Enforcement" will be integrated in the two FOCETA use case demonstrators as part of the work package led by Son. What are the main challenges and next steps towards this integration?

A6 / Son: The first challenge on my side is to understand the development done by the FOCETA partners. There are so many tools have been developed, using different methods and domain background ranging from control to machine learning and formal methods. The second challenge is to know how they can be integrated into the common demonstrators. We aim to gradually increase the complexity

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

and readiness level of the use cases, hence more collaboration efforts in the later phase of the project. The next step is to get the demonstration ready with extensive validation in simulation environment. In the third project year [i.e., between October 2022 and September 2023], we plan to further enhance the validation efforts with physical system integration.

A6 / Bettina: In the first two project years, we focused on developing our methods. As explained by Son, now the challenge is to integrate them in the FOCETA use cases and to assist the industrial partners in doing so. One of the main challenges of this step will be the high complexity that the real-world FOCETA case studies bring with them. The propotype tools that the academic partners developed in the past two years are mostly tested on simpler benchmarks. Especialy for the shield synthesis tool that we developed at Graz University of Technology, I am very curious to see how our shields are applied in a real autonomous driving senario. For the integration, we will from now on work closely together with the industrial partners to understand how our shields can be best integrated into the existing system implementations.

Q7: The FOCETA consortium gathers many European and international partners, including leading academica and research centers, as well as prominent industrial partners. What are the main challenges related to the collaboration between such different organisations? Could you give any examples of key differences in the approach of academic/research centers vs. industrial partners in the research activities within FOCETA? Are these differences perceived as a constraint or do they act as drivers towards disruptive solutions to tackle the complex issue of Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems?

A7 / Bettina: The FOCETA consortium consists of excellent academic and industrial partners, and I am still very exited about the opportunity to work with them. We had never collaborated with most of the FOCETA partners before the project start. Therefore, I am very happy to have expanded my research network and to have learned a lot from all the partners involved. All industrial partners involved in both use cases gave workshops, seminars and demos, and worked closely with academic partners to integrate the monitoring tools, which helped us gaining a much clearer vision about safety and what should be monitored. Personally, I perceived the medical use case as a challenge in the beginning, since I lacked any knowledge about anesthesia or drug infusion pumps, and I was therefore uncertain on how the methods developed by the academic partners could fit into this use case. R G B MEDICAL DEVICES SA transferred a lot of knowledge to all partners involved and gave several workshops, but still in the beginning the progress we made for this use case felt slow. After two years, it is now nice to see that there are actually many similarities between the two FOCETA use cases and in how we adress safety in them.

A7/ Son: That is an interesting question. Indeed, it is great for us to have the chance to work together with such excellent partners in FOCETA. We have learnt much from these collaborations. It is also nice that we are working on common use case applications, where the technologies are consolidated into a similar domain. There are a few minor challenges which reflect the different academic and industry working environment. For example, the academic partners often prefer state-of-art and advanced algorithms, while the industrial partners rather put the focus on the maturity of the tool, its robustness, as well as testing and industrial standard (safety, requirements, middleware, etc.). The academic colleagues also have more flexibility in exploring different methods and techniques.

These differences were challenging in the beginning of the project, but in my view, we have reached a good common ground now, with a better level of understanding from both sides. Through collaboration, we have been exploiting and complementing each other's strengths, combining them into disruptive solutions and showing them via concrete use case demonstrations.

Q8: What is the impact of the FOCETA project on your field of research? How will the Graz University of Technology and Siemens Digital Industries benefit from the innovations developed within FOCETA?

A8 / Bettina: To automatically construct correct-by-construction shields is my main research topic since several years. Nevertheless, I only applied my shields on textbook examples and several video games so far, always with the vision to have them one day integrated in a self-driving car. FOCETA gave me the opportunity to integrate shields to enforece safety in a real-world valet parking senario, which is very exiting for me.

A8/ Son: FOCETA project helps us to accelerate our R&D activities on multiple strategic topics, from modelling to control and AI system developent, as well as safety requirements and scenario generation. The FOCETA use cases development allows us to obtain extra experience on automated valet parking and depth of anaesthesia test platforms from the system level point of view. Siemens Digital Industries has obtained quite some interesting results which have been presented both extenally to internation conference and internally to our customers.

Q9: Bettina, it is a matter of fact that women are underrepresented in the fields of engineering and research on artificial intelligence. Based on your experience and career, what could help progress towards a more gender-balanced occupation in this field?

A9 / Bettina: I think that one of the problems is that students often only see the final product, which capabilities are often very impressive and give the impression that one needs to be all-knowing to be able to contribute to the field. This might in particular discourage female students. We need to show high-school as well as university students that this is a very young field, where there is a lot left to be done, and we need many people with many different talents, interests and skills to make advances in the field. We need people that have a strong mathematical background to advance current algorithms, and we need people interested on the practical aspects to train and to test the networks. Furthermore, we need people interested in the psychological and sociological aspects of human-AI-interaction to study on what is needed to establish trust in the decisions made by AI-based systems. Therefore, I think that to attract more female researchers in the field we need to show them (1) that the field is highly interesting from many different aspects and (2) give them the confidence that they can contribute to the field.

DISCLAIMER - The information, statements and opinions in the above interview are personal views of the individuals involved in the FOCETA project and do not necessarily reflect the views of the FOCETA consortium as a whole, nor of the European Commission. None of them shall be liable for any use that may be made of the information contained herein.

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

GET TOGETHER

In this issue of our Newsletter you will find one feedback article from the H2020 Summer School co-organised by FOCETA, as well as a selection of major upcoming conferences which are of interest for the FOCETA community.

FEEDBACK ARTICLE ON THE JOINT H2020 OPENDR - FOCETA SUMMER SCHOOL: CONTINUOUS ENGINEERING AND DEEP LEARNING FOR TRUSTWORTHY AUTONOMOUS SYSTEMS, 3RD-7TH OCTOBER 2022, THESSALONIKI, GREECE

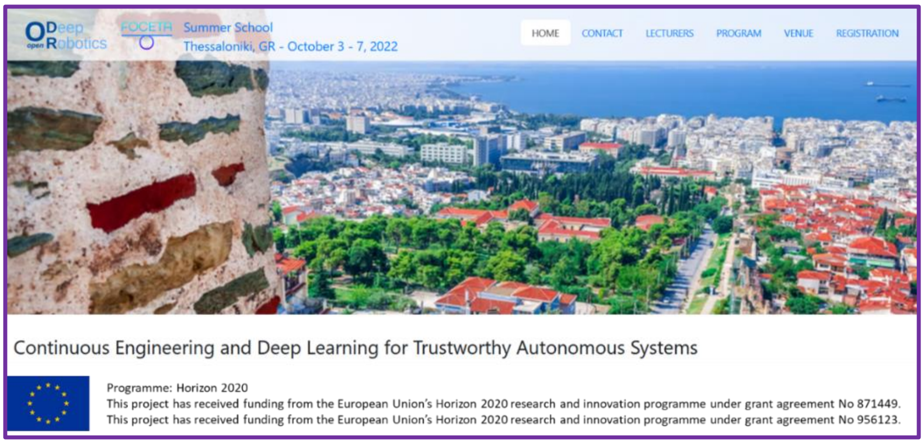

Figure 8 : Homepage of the H2020 OpenDR-FOCETA Summer School website.

The H2020 OpenDR-FOCETA Summer School "Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems" took place between 3rd-7th October 2022 in Thessaloniki, Greece, hosted by the Aristotle University of Thessaloniki. The summer school was jointly organised by the Horizon 2020 sister projects FOCETA and OpenDR (Grant Agreement n°871449) as part of the training activities planned in the framework of each of the projects.

The aim of the Summer School was to demonstrate the approaches and methodologies implemented in both H2020 projects to cope with the complex issues related to Continuous Engineering and Deep Learning for Trustworthy Autonomous Systems. As such, the event was specifically designed for graduate and doctoral students as well as young researchers.

The summer school offered a dense programme with invited talks by distinguished researchers and industrial experts, as well as talks by the project participants that reflected the achievements in the two ongoing projects. Two days were dedicated to presentations by the OpenDR project partners and two other days to talks by the FOCETA partners. The last day was dedicated to hands-on and lab sessions on demonstrating tools related to the subjects of the FOCETA and the OpenDR projects.

This successful event recorded 75 registerations and over 60 attendees, and led to several interesting discussions.

Several social activities had also been organised over the week, which enhanced the networking between the attendants while discovering the beauties of the city of Thessaloniki.

The summer school presentations will be available for download soon on the FOCETA website. Stay tuned by following us on FOCETA Linkedin.

UPCOMING DISSEMINATION EVENTS

NEURIPS 2022 - 36TH CONFERENCE ON NEURAL INFORMATION PROCESSING SYSTEMS, 28TH NOVEMBER-9TH DECEMBER 2022, NEW ORLEANS, LOUISIANA, US

NeurIPS 2022 brings together a broad community around machine learning, artificial intelligence, and neural information processing. **FOCETA partner University of Liverpool will give a presentation** based on their paper entitled "Recursive Reinforcement Learning".

Source: https://neurips.cc/Conferences/2022/

FSTTCS 2022 - 42ND IARCS ANNUAL CONFERENCE ON FOUNDATIONS OF SOFTWARE TECHNOLOGY AND THEORETICAL COMPUTER SCIENCE, 18-20TH DECEMBER 2022, CHENNAI, TAMIL NADU, INDIA

FSTTCS 2022 is the 42nd conference on Foundations of Software Technology and Theoretical Computer Science. It is organised by IARCS, the Indian Association for Research in Computing Science, in association with ACM India. It is a forum for presenting original results in foundational aspects of Computer Science and Software Technology. **FOCETA partner University of Liverpool will give a presentation** based on their paper entitled "Natural Colors of Infinite Words".

Source: https://www.fsttcs.org.in/2022/

This project has received funding from European Union's Horizon 2020 research and innovation programme under grant agreement N°956123.

FM 2023 - 25TH INTERNATIONAL SYMPOSIUM ON FORMAL METHODS, 6-10TH MARCH 2023, LÜBECK, GERMANY

FM 2023 is the 25th international symposium in a series organized by Formal Methods Europe (FME). The topics covered include the development and application of formal methods in a wide range of domains including software, cyber-physical systems and integrated computer-based systems. In case of acceptance of their papers, **FOCETA partners AIT Austrian Institute of Technology and Graz University of Technology plan to give a presentation** each at FM 2023.

Source: https://fm2023.isp.uni-luebeck.de/

FOSSACS 2023 - 26TH INTERNATIONAL CONFERENCE ON FOUNDATIONS OF SOFTWARE SCIENCE AND COMPUTATION STRUCTURES, 22-24TH APRIL 2023, PARIS, FRANCE

FoSSaCS seeks original papers on foundational research with a clear significance for software science. The conference invites submissions on theories and methods to support the analysis, integration, synthesis, transformation, and verification of programs and software systems. In case of acceptance of their paper, **FOCETA partner University of Liverpool plans to give a presentation** at FoSSaCS 2023.

Source: https://etaps.org/2023/fossacs

ICCPS 2023 - 14TH ACM/IEEE INTERNATIONAL CONFERENCE ON CYBER-PHYSICAL SYSTEMS, 9-12TH MAY 2023, SAN ANTONIO, TEXAS, US

ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS) is the premier forum for presenting and discussing the most significant recent technical research contributions in the field of CPS. ICCPS will be part of the CPS-IoT Week 2023. In case of acceptance of their paper, **FOCETA partners AIT Austrian Institute of Technology plans to give a presentation** at ICCPS 2023.

Source: https://iccps.acm.org/2023/

ICRA 2023 - 2023 IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION, 29TH MAY-2ND JUNE 2023, LONDON, UK

2023 IEEE International Conference on Robotics and Automation (ICRA) is one of the most prestigious events in Robotics and Automation, bringing together the world's top academics, researchers, and industry representatives. Many FOCETA partners have applied to the conference: in case of acceptance of their papers, **FOCETA partners University of Liverpool, DENSO AUTOMOTIVE DEUTSCHLAND GmbH** and **AIT Austrian Institute of Technology plan to give a presentation at ICRA 2023.**

Source: https://www.icra2023.org/

Follow the FOCETA project on LinkedIn

This project has received funding from European Union’s Horizon 2020 research and innovation programme under grant agreement N°956123.