Transactions of the SDPS:
Journal of Integrated Design and Process Science
25 (2), 2021, 27-47
DOI 10.3233/JID210010
http://www.sdpsnet.org

# A Research Agenda for Smarter Cyber-Physical Systems

Danny Weyns[a,b*], Jesper Andersson[b], Mauro Caporuscio[b], Francesco Flammini[b,c], Andreas Kerren[b,d], and Welf Löwe[b,e]

[a] *Department of Computer Science, Katholieke Universiteit Leuven, Belgium*
[b] *Department of Computer Science and Media Technology, Faculty of Technology, Linnaeus University, Sweden*
[c] *School of Innovation, Design and Engineering, Division of Product Realisation, Mälardalen University, Eskilstuna, Sweden*
[d] *Department of Science and Technology, Linköping University, Sweden*
[e] *Softwerk AB, Sweden*

**Abstract** With the advancing digitisation of society and industry we observe a progressing blending of computational, physical, and social processes. The trustworthiness and sustainability of these systems will be vital for our society. However, engineering modern computing systems is complex as they have to: i) operate in uncertain and continuously changing environments, ii) deal with huge amounts of data, and iii) require seamless interaction with human operators. To that end, we argue that both systems and the way we engineer them must become *smarter*. With smarter we mean that systems and engineering processes adapt and evolve themselves through a perpetual process that continuously improves their capabilities and utility to deal with the uncertainties and amounts of data they face. We highlight key engineering areas: cyber-physical systems, self-adaptation, data-driven technologies, and visual analytics, and outline key challenges in each of them. From this, we propose a research agenda for the years to come.

**Keywords**: Smarter systems, trustworthiness, sustainability, cyber-physical systems, self-adaptation

## 1. Introduction

The advancing digitisation of society and industry leads to an increasing blend of computational and physical processes (Baheti & Gill, 2019; Lee, 2008). This progressing integration of cyber and physical elements combined with a seamless integration of social elements (Liu, Yang, Wen, Zhang, & Mao, 2011; Zeng, Yang, Lin, Ning, & Ma, 2020) and the increasing amount of data that needs to be processed results in computing systems with software in a predominant role. Consequently, virtually everything we do today relies directly or indirectly on software. The future of our society depends on the trustworthiness of these systems, i.e., the compliance of the systems with their business, technical, and legal requirements, and their sustainability, i.e., the longevity of these systems and their infrastructure. Examples are reliable and continual intelligent traffic control (Lin, Wang, & Ma, 2017), smart grids (Tuballa & Abundo, 2016), and

*Corresponding author. Email: danny.weyns@kuleuven.be. Tel: (+32)474-208251.

1092-0617/©2021 - The authors. Published by IOS Press
This article is published online with Open Access and distributed under the terms of the Creative Commons Attribution Non-Commercial License (CC BY-NC 4.0).

manufacturing automation in Industry 4.0 (Zhou, Liu, & Zhou, 2015). Since these systems operate in uncertain and continuously evolving environments, we argue that both the systems and the way we engineer them must become smarter. We define "smarter" as follows:

> *Systems and engineering processes continuously adapt and evolve themselves from experience and stakeholder input through a perpetual process that continuously improves their capabilities and utility to deal with the uncertainties and new data they face throughout their lifetime.*

As such, smarter is a relative concept that expresses an increase of capabilities and utility of a system over its lifetime. These system enhancements are obtained by mitigating uncertainties and processing new data that the system encounters over time. Increasing capabilities can be measured by enhancements in the functional capacity and abilities of the system (i.e., the system can deal with tasks that it was not able to deal with before). Increasing utility can be measured by enhancements in the qualitative concerns of the system (i.e., the system can perform tasks more efficiently, reliably, etc.).

To tackle the challenges of systems that blend cyber, physical, and social elements, other researchers have argued for the need of smartness. We discuss a selection of relevant work. Jazdi (2014) highlights the need to equip Industry 4.0 systems with smart actuators, sensors, and telecommunication technologies, providing these systems access to the higher-level processes and services. Bures et al. (2015) emphasise that smartness of computing system enables them to deal with environment dynamics and uncertainty, cope with external threats, and optimise their behaviour to achieve the best possible outcome. The authors highlight that smartness is primarily implemented in software typically through cooperative behaviour, self-awareness, self-adaptation, and self-optimisation. Yu and Xue (2016) refer to smartness of the electricity grid as the integration of enabling information and communication technology with other advanced technologies that enable electric energy generation, transmission, distribution, and usage to be more efficient, effective, economical, and environmentally sustainable. Koutsoukos et al. (2018) investigate smart transportation systems using a modelling and simulation environment. Smartness in this context relates to the ability of the system to deal with attacker–defender behaviour, including vulnerability analysis to traffic signal tampering, resilient sensor selection for forecasting traffic flow, and resilient traffic signal control in the presence of denial-of-service attacks. Tavčar and Horváth (2019) survey smart cyber-physical computing, distinguish four levels of smartness mapping to increasing challenging types of changes to be tackled by the systems. The article outlines a set of techniques to equip computing system with smartness, relying on reasoning, learning, adapting, and evolving capabilities. Finally, Zeng et al. (2020) emphasises the role of smart spaces in cyber-physical-social systems with typical applications such as smart home, smart transportation system, and smart medical service system. The authors highlight the importance of understanding user intentions as a critical aspect regarding smartness.

The need for smart systems is also underpinned by recent discussions within the international research community, see e.g., Bures et al. (2018); CPSWeek (2021), reflected in several surveys, see e.g., Gunes et al. (2014); Stankovic (2016); Zhou et al. (2015), and recent funding programs, see e.g., EU (2021b); NSF (2021). The required shift relates to how the next-generation computing systems should be developed and what methods and techniques need to be brought together in order to achieve the required smartness while dealing with the complexity arising from the scale, connectivity, and inherent uncertainty of these systems. The inherent blend of cyber, physical and social components demands for novel engineering approaches as these systems present a combination of characteristics that existing modelling and development methods have difficulties addressing, including a widespread uncertainty and disruptions in a broad variety of changing contexts and continuous encounter with new situations, humans as integral elements, and the presence of large volumes of data that needs to be processed.

Our work complements these existing perspectives on smartness of systems by taking a more holistic perspective that integrates systems operation with the processes to engineer them. In particular, we intentionally use the term "smarter" underpinning that adaptation and evolution is an enduring process that concerns systems operation and the engineering processes, spanning from inception of systems to and throughout their operation in the real world. Key characteristics of smarter systems and processes are *adaptation* that refers to the ability of mitigating anticipated uncertainty in order to keep satisfying the goals (Esfahani & Malek, 2013; Mahdavi-Hezavehi, Avgeriou, & Weyns, 2017), whereas *evolution* refers to the ability of accommodating unanticipated uncertainty in order to handle goal changes and novelty (Calinescu, Mirandola, Perez-Palacin, & Weyns, 2020; Weyns, Caporuscio, Vogel, & Kurti, 2015).

Our perspective on smartness relies on the observation that modern computing systems are often feedback-driven characterised by a tight coupling between software and physical elements, the presence of large volumes of data, and a seamless integration of humans in the loop (Selic, 2020; Sztipanovits et al., 2012; Zeng et al., 2020). Since these systems face uncertainties that are difficult or even impossible to predict before deployment, engineers may not be able to obtain sufficient knowledge to make all design decisions before the system is deployed. Hence, some design decisions are postponed until after deployment, that are then enacted through continuous adaptation and evolution (Weyns, 2020). In this view, system engineering and system operation get blended (Baresi & Ghezzi, 2010). The ubiquity and scale of systems and the rapid integration with of data-intensive resources such as smart cameras introduce the need for handling previously unknown amounts of data. This requires novel advanced data driven technologies that enable smarter systems to learn from experiences and examples by exploiting data (Gandomi & Haider, 2015). Smarter systems and their engineering processes rely on knowledge obtained from stakeholders to drive system adaptation and guide its evolution. This calls for advanced technologies for incorporating humans in the loop. Visual analytics (Thomas & Cook, 2005) is the field that investigates ways to better comprehend large and complex data by combining the strengths of human and computational data processing, providing the means for comprehensible interaction between systems and stakeholders. The seamless integration of continuous adaptation and evolution supported by data driven technologies and visual analytics aims at mitigating the effects of uncertainty and providing means for satisfying stakeholder requirements in a trustworthy and sustainable manner across the lifetime of the system. This seamless integration is the basis of being smarter as posited in this paper.

The aim of this paper is to outline a research agenda for the engineering and operation of smarter systems. Our particular focus is on computing systems that blend cyber, physical, and social elements. Figure 1 gives an overview of the approach we follow to devise this research agenda. We leverage on the state of the art in smart systems to define key objectives for smarter systems. Driven by these objectives we identify challenges in four key areas of smarter systems: cyber-physical systems that provide a set of basis challenges, and self-adaptation that targets uncertainties, data-driven technologies that target large amounts of data, and visual analytics that targets humans in the loop. With this analysis in hand, we outline a research agenda for smarter cyber-physical systems comprising three themes: assurances for unknowns, self-explainability, and smarter ecosystems.

It is important to highlight that we selected three key emerging demands for smarter systems based on the analysis of the state of the art (uncertainties, large amounts of data, and humans in the loop). For each of these demands we defined concrete challenges in corresponding domains (cyber-physical systems as a basis, self-adaptation, data-driven technologies, and visual analytics) resulting in three research themes (assurances for unknowns, self-explainability, and smarter ecosystems). As such, the research agenda takes a particular focus. Other complementary challenges exist, ranging from energy efficiency (Schmidt & Åhlund, 2018) up to para-functionalities such as empathy and emotions that are extensively studied in socialised robotics or embodied AI (Kephart et al., 2019).

This paper targets primarily researchers with an interest in smarter cyber-physical systems, although all stakeholders with an interest in the development and operation of smarter computing systems may find inspiration in the work presented in this paper. In Section 2, we start with outlining the key objectives of smarter systems both for systems and engineering processes. Driven by these objectives, we then highlight

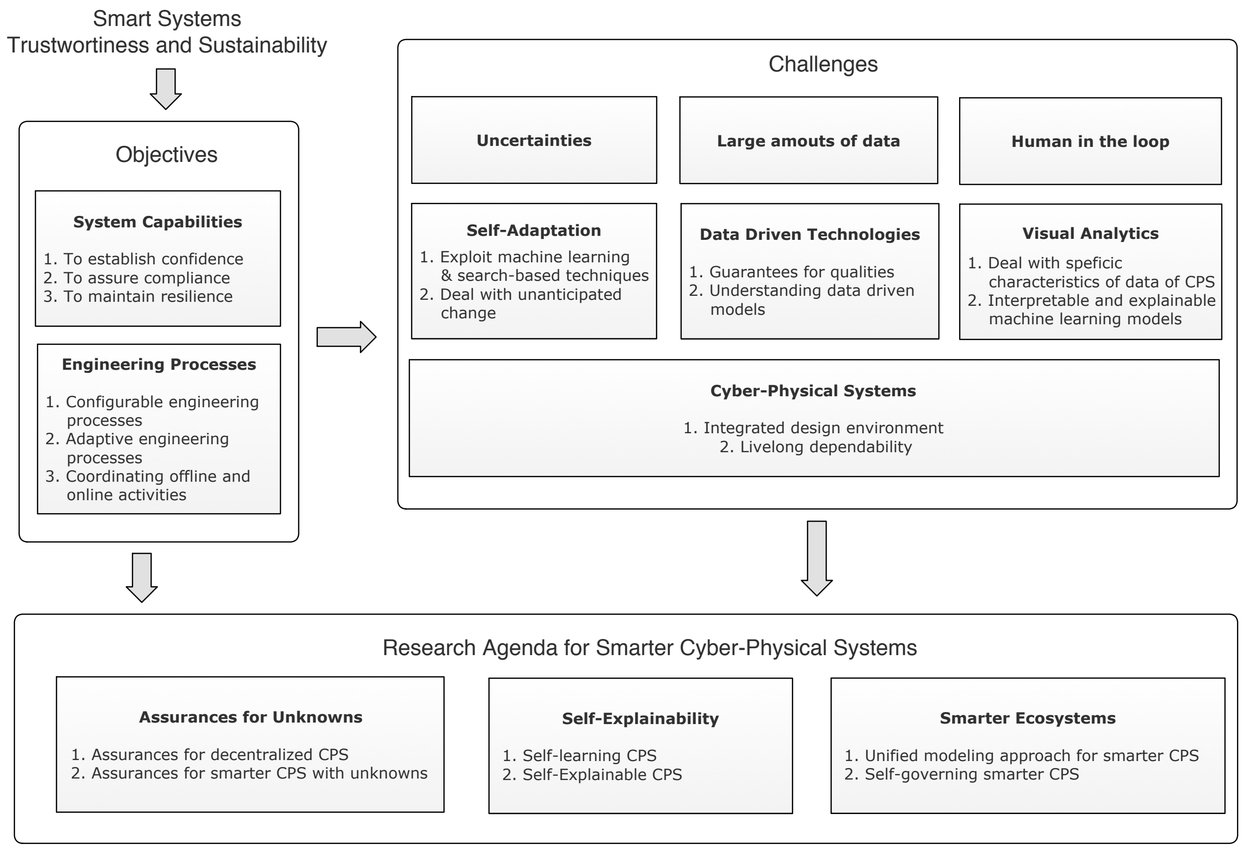


**Fig. 1. From Objectives to Challenges to the Research Agenda**

four core areas that underlie the engineering and operation of smarter systems and highlight key challenges in each of them in Section 3. Section 4 then brings together the pieces, outlining a research agenda to tackle the challenges. Finally, we wrap up and draw conclusions in Section 5.

## 2. Key Objectives for Smarter Systems

The key research objectives for smarter systems centre on two drivers that are broadly considered a as key for future computing systems that operate under uncertainty: trustworthiness and sustainability, see for instance EU (2021a). Trustworthiness refers to confidence of stakeholders in smarter systems and the compliance of systems with business, technical, and technological requirements, as well as legal obligations. Gol Mohammadi (2019) argues that trustworthiness is a key success factor in the acceptance and adoption of cyber-physical systems. Establishing trust requires methods that cover all phases of development and operation: requirements engineering, system design, reliable evaluation, run-time maintenance, and evidence-based assurance. We share this point of view in this work. Sustainability generally refers to the *endurance* of systems and processes. In particular, technical sustainability refers to "the longevity of information, systems, and infrastructure and their adequate evolution with changing surrounding conditions" (Becker et al., 2015). We share that view in this work. To achieve the drivers trustworthiness and sustainability, we look at smarter systems from two complementary angles: *system capabilities* and *engineering processes*.

### 2.1. System capabilities objectives

To ensure trustworthiness and sustainability, we distinguish three core capabilities of smarter systems, leveraging on the design principles proposed in (Gupta et al., 2011; Tavčar & Horváth, 2019): *confidence*, i.e., stakeholders are confident that the systems can be relied upon and will remain doing so over time; *compliance*, i.e., the systems behave according to specification throughout their lifetime; and *resilience*, i.e., the systems provide an acceptable quality of service also when facing instability and uncertainty inherent to ever-changing settings. Hence, the first key objective is to establish and maintain confidence, compliance, and resilience throughout the lifetime of smarter systems that operate under continuous change in stakeholders' goals, environments, and the systems themselves. In more detail:

(1). **To establish confidence:** establishing confidence among the system' stakeholders requires the designers of smarter system to formulate and translate their concerns to system requirements and constraints. This raises several open questions: Can we build smarter systems that are trustworthy and sustainable by construction, and if not, how can we ensure trustworthiness and sustainability during operation? How can we make the traditional black-box analysis methods transparent for the analysts and other stakeholders? How can we trust and sustain a smarter system that is composed of a variety of heterogeneous components (sub-systems and stakeholders)?

(2). **To assure compliance:** assuring that the behaviour of smarter systems stay within the business, technical, and legal constraints requires designers to formulate compliance requirements and techniques to assure them throughout the lifetime of systems. This raises open questions, such as: How to assure system compliance in an ever-changing and uncertain setting and how can we test and verify compliance requirements? How to employ data science and analytics techniques to provide guidelines and supporting strategies—well-justified in the collected data—to sustain the system in the long term?

(3). **To maintain resilience:** maintaining an acceptable level of service of a operating smarter system in the face of changes and faults requires appropriate methods that span the full lifetime of a system. How can we build smarter systems that adjust themselves in an ever-changing and uncertain environment without losing trust, compliance and quality? How to integrate adaptation processes with evolution processes of smarter systems to satisfy short- and long-term stakeholder concerns in a continuously evolving operational context?

### 2.2. Engineering processes objectives

Ensuring trustworthiness and sustainability requires advanced engineering processes that span the lifetime of smarter systems, from development to operation and maintenance. Leveraging on insights proposed in (Andersson et al., 2013; Tavčar & Horváth, 2019), we argue that smarter systems require engineering processes that tightly integrate development time activities and runtime activities, uniting evolution and adaptation. This implies that a deployed smarter system is equipped with mechanisms to identify the need for change and coordinate with offline development support. Hence, the second key objective is to devise new principled engineering processes for smarter systems that seamlessly integrate engineering activities with system activities, spanning across the full system life-cycle. In more detail:

(1). **Configurable engineering processes:** smarter systems require configurable life-cycle processes with the necessary variability to be adapted and evolved as systems and operating conditions change. This raises open questions such as: What are the requirements of configurable engineering processes and how can they be defined? How can such processes be verified for completeness and correctness properties, and to what extent is this needed? What types of process models are suitable for smarter systems?

(2). **Adaptive engineering processes:** processes for engineering smart systems need to adapt and evolve in support of trustworthy and sustainable system adaptations and evolution. This raises open questions as: What type of offline and online mechanisms for engineering processes are required to perpetually support trustworthiness and sustainability of smarter systems? What are the triggers for dynamic reconfiguration of engineering processes? How to provide on the fly guarantees for completeness and correctness properties of processes?

(3). **Coordinating offline and online activities:** seamless integration of evolution and adaptation requires the coordination of online (machine-driven, human-supported) and offline (human-driven, tool-supported) activities. This raises challenging questions such as: How may activities of running systems share data and knowledge with the offline activities? What role may simulation play in unifying offline and online activities? What type of abstractions and coordination mechanisms connect cyber, physical, and human elements, within and across development, adaptation, and evolution activities?

## 2.3. Illustrative industrial scenario

We illustrate how we may engineer new smarter systems using a scenario of a smart grid. Parts of the scenario are intentionally speculative and discuss future capabilities where research contributes new knowledge that drive the development of methods, techniques, and tools for engineering and operating smarter systems.

*Smart Grid*. The power grid comprises power providers (the generation side) and consumers (the consumption side) that are connected through transmission and distribution lines. The grid is operated by one or more control centres. Trends are pushing control closer to the equipment and adding capabilities to react autonomously to events, without human intervention. The power grid domain is going through a paradigm shift due to multi-fold challenges.

*Challenges*. On the power generation side, the rapid uptake of solar panels and other forms of local energy generation systems, energy production is no longer owned by the traditional large players alone. At the same time, there is an ecological (and political) drive to produce clean energy to curb the greenhouse gas emissions coming from the fossil-based energy sources. The clean energy production from renewable sources is intermittent and, hence, comes with uncertainty in both the amount of energy produced and the stability of the produced power in terms of magnitude, phase, and frequency. On the consumption side, the traditional bulk of industrial consumers are also shifting. A classic example is the energy consumed by the data centres, which is for instance slated to be 20% of the total energy consumption in the US. Similarly, millions of plug-in electric vehicles will emerge in the coming years that will disrupt the charging infrastructure demands. Rapid urbanisation in different parts of the world is generating a modified need for large commercial establishments and bulk residential consumer base. Importantly, the nature of the load in the setups mentioned above is mostly non-linear, stressing the regular operation of the grid. In order to meet the challenges, operators and grid owners have to find new solutions that are *trustworthy* and *sustainable*, i.e., ensure the compliance of the grid with its business, technical, and legal requirements, and its longevity regardless of the changing conditions it will face throughout its lifetime.

*A Smarter Grid*. The multitude of challenges requires a holistic solution that considers the power system as an entailed cyber-physical system. The resulting smarter grid offers interfaces for monitoring and operating the grid by software. A typical configuration combines smart meters and smart appliances on the consumption side and renewable energy resources and smart distribution equipment on the generation side. Evidently, the smarter grid needs to be operational 24/7. Any change action in the grid, from a simple adaptation of the running configuration up to an invasive upgrade of its functionality, needs to be applied life, without any downtime. As the smart grid needs to provide its services without interruption over a long period of time, change management necessarily becomes a perpetual process. Uncertainty is impacting smart grid design and operations. The complexity and decentralisation of the system make it difficult to comprehend. Uncertainty is caused by events that are impossible to anticipate in time, for example:

(1). Blackouts caused by natural disasters or cascading human errors;

(2). Intentional cyber-warfare attacks by hackers;

(3). Unprecedented grid pollution due to unaccountable generation and load that fails to meet peak demand, leading to brownouts, and short interruptions.

One trend for grid services is shorter response times and planning horizons. Another trend is to use streams of real-time data in decision-making. Seamless operation of such services with minimal disruption and downtime calls for interdisciplinary research on autonomous adaptation and evolution to mitigate uncertainty. Health care, industries, charging stations, and many ancillary services rely on the electricity grid. *Establishing confidence* among these stakeholders and *assuring compliance* with their requirements is crucial for the trustworthiness and sustainability of the smarter grid. Such complex systems are not only vulnerable to potential faults, but also malicious attacks if the cyber-infrastructure is not maintained proactively. This calls for self-protection against cyber-threats that monitors and analyses the system to detect malicious behaviours, and plan and enact adaptations to protect the system and *maintain its resilience* at any time.

A smarter grid will not be completely autonomous and human operators and engineers will play an important role; i.e., the synergy between the operating system and its stakeholders will be crucial in making the grid smarter. The grid domain is subject to continuous change, with new technologies emerging virtually every day; consider for instance new emerging methods for large-scale energy storage. This calls for *configurable engineering processes* that seamlessly align their activities with the changing technologies and operating conditions the grid faces throughout its lifetime. *Adaptive engineering processes* enable dynamically adjustment of engineering activities to handle uncertainties caused by incidents as listed above and evolve the grid with new emerging technologies. Ensuring resilient operation of the grid requires *coordinating the activities* of engineers supported by tools with the activities of the operational grid. Support for bi-directional comprehensive communication between the system and stakeholders will therefor be crucial.

## 3. Core Areas and Challenges

In this section, we summarise four core research areas that are central in the engineering of smarter systems: cyber-physical systems, self-adaptation, data driven technology, and visual analytics. Cyber-physical systems provide a basis for smarter systems. The other three areas target key emerging demands for smarter systems: self-adaptation targets uncertainties, data-driven technologies target large amounts of data, and visual analytics are key to humans in the loop. We highlight representative state of research in each of the four areas and outline key research challenges for each area. Then we explain how the four areas complement one another as a basis for a research agenda.

### 3.1. Cyber-physical systems

Cyber-Physical Systems (CPS) are engineered systems that are built from, and depend upon, the seamless integration of computational and physical components (Bures et al., 2015; Lee, Bagheri, & Kao, 2015). Cyber-physical systems are becoming increasingly complex, critical, ubiquitous and pervasive. Research shows that the complexity is a result of three main factors (Banerjee et al., 2012; Sztipanovits et al., 2012; Tavčar & Horváth, 2019; Tokody, Papp, Iantovics, & Flammini, 2019): (i) size of the software and of the whole system (system-of-systems) due to non-straightforward functional requirements to be fulfilled; (ii) hardware and software heterogeneity due to diverse embedded systems architectures, protocols, manufacturers and connection facilities, possibly including legacy devices; and (iii) distribution due to large networks of connected devices, including the Internet of Things (IoT), Industry 4.0, and domains with strict goals such as Intelligent Transportation Systems and e-health. These complexities raise challenges; we highlight two of the key challenges that are driven by the objectives of smarter systems.

A first challenge is to devise *integrated design environments* for Model Driven Engineering (MDE) of cyber-physical systems leveraging on standard high-level languages like the Unified Modelling Language (UML) and the Systems Modelling Language (SysML) (D'Angelo, Napolitano, & Caporuscio, 2018). A crucial advanced feature of an integrated design environment for smarter cyber-physical systems will be the ability to handle the heterogeneity that arises from compositional arrangements of analogue and digital hardware, control and application software and middleware, and coded and run-time obtained cyberware. Furthermore, multi-paradigm modelling —including multiple levels of abstraction, multi-formalism modelling, and meta-modelling (Ciccozzi, Tichy, Vangheluwe, & Weyns, 2019; Fitz, Theiler, & Smarsly, 2019)— with domain specific languages and appropriate model-to-model transformations would enable a set of automated analysis supporting cyber-physical system assessment and certification against international security and safety standards (Flammini, Marrone, Nardone, Caporuscio, & D'Angelo, 2020), e.g., ISO/IEC 15408, IEC 61508. Tackling the first challenge will contribute to the trustworthiness and sustainability of smarter cyber-physical systems by establishing confidence among stakeholders and ensuring compliance of cyber-physical systems with their requirements, and support configurable and adaptive engineering processes.

A second challenge is providing the necessary levels of *livelong dependability* in the face of uncertainty, including reliability, security and performance of CPS (Bennaceur et al., 2019; Pagliari, Mirandola, & Trubiani, 2020; Ratasich et al., 2019). Cyber-physical systems often operate for many years during which these systems are exposed to vulnerabilities unknown at development time. Software upgrades are not straightforward in cyber-physical systems due to potential dependability implications, e.g., constraints of safety certification. Hence, securing critical and non-critical functionality needs to be distinguished. Besides cyber-security, dependability issues need to be addressed, such as protocol incompatibilities between heterogeneous devices manufactured by diverse vendors. Understanding and mitigating such causes is crucial for cyber-physical systems (Caporuscio, Flammini, Khakpour, Singh, & Thornadtsson, 2020). Featuring higher levels of autonomy and intelligent behaviour can fuel advanced prediction realising "proactive dependability." New paradigms like "digital twins" support real-time prediction of problems through run-time "what if" simulations (Tao, Qi, Wang, & Nee, 2019). Tackling the second challenge will contribute to the trustworthiness and sustainability of smarter cyber-physical systems through the resilience of the systems in face of changes and the coordination of offline and online activities to achieve a seamless integration of evolution and adaptation.

## 3.2. Self-adaptation

In 2003, IBM released a manifesto referring to "a looming software complexity crisis" that was caused by the increasing complexity of installing, configuring, tuning, and maintaining computing systems (IBM, 2003; Kephart & Chess, 2003). This led to the notion of "self-adaptation," i.e., systems that can adapt themselves autonomously or with minimal human intervention. The motivation for self-adaptation is dual; on the one hand it offers a means to free system administrators from the details of managing computing systems that run 24/7; on the other hand it enables systems to deal with uncertainties that were difficult to foresee before deployment (Weyns, 2020). A common approach to realise self-adaptation is by means of an external feedback loop that realises four basic functions: Monitor, Analyse, Plan, and Execute. These functions share Knowledge, hence, the model is often referred to as MAPE-K. Researchers have argued for an architecture perspective on engineering self-adaptive software systems (Garlan, Cheng, Huang, Schmerl, & Steenkiste, 2004; Kramer & Magee, 2007), providing generality of concepts and an appropriate level of abstraction to define self-adaptive systems and to reason about adaptation at runtime. Blair, Bencomo, and France (2009) emphasise the role of software models at runtime as an extension of model driven engineering techniques to the runtime context. Runtime models provide abstractions of the system and its goals serving as a driver and enabler for automatic reasoning about system adaptations during operation. A recent analysis concludes that the area is currently in the phases of internal and external enhancement and exploration (Weyns, 2019). This implies that research is consolidating by showing initial evidence of its value, but industrial validation is needed. We highlight two open challenges in this area that are driven by the objectives of smarter systems.

A first important challenge is *exploiting machine learning and search-based techniques.* that can play a central role in virtually every stage of adaptation, from processing large amounts of data, performing smart analysis and machine-man co-decision making, to coordinating adaptations in large-scale decentralised settings (Gheibi, Weyns, & Quin, 2021). First results in this direction have been already been presented, see for instance Cheng, Ramirez, and McKinley (2013); Van Der Donckt, Weyns, Quin, Van Der Donckt, and Michiels (2020). Smarter systems inherently will require techniques that are efficient as they need to be applied at runtime. Tackling the first challenge will primarily contribute to trustworthiness of smarter systems by assuring their compliance with stakeholder requirements and maintaining resilience of the operating system during system adaptation.

A second challenge is dealing with *unanticipated change*, a challenging characteristic of smarter systems. An intriguing question is to what extent software can handle conditions that were not anticipated when developed. One perspective on tackling this problem is to seamlessly integrate self-adaptation, i.e., machine-driven adaptation to deal with known unknowns, with evolution, i.e., the human-driven updates to deal with unknown unknowns, which goes back to the vision proposed in the seminal work of Oreizy et al. (1999). ActivFORMS is one approach in this direction (Weyns & Iftikhar, 2019). Another perspective would be to conceive a system as a dynamic composition of learning processes, and enhance the system then with self-learning capabilities. Calinescu et al. (2020) present an interesting survey on the perception of the research community on handling unanticipated changes. Tackling the second challenge will be pivotal to a seamless integration of online machine-driven activities with offline human-driven activities, contributing to the trustworthiness and sustainability of smarter systems.

## 3.3. Data driven technologies

The ubiquity of environment interaction, computing, communication, and storage technologies provides access to previously unknown amounts of data. The objective of data driven technologies is to learn from experiences and examples by exploiting data. To that end, these technologies transform data into information into actionable knowledge (value) while managing challenging quantities (volume, velocity) and qualities (variety, veracity, validity) of data (Gandomi & Haider, 2015). Dealing with data quantity and quality relies on technologies such as parallel and real-time computing and compiler technologies (including meta-modelling, interpreting, and composition of heterogeneous data sources), while transforming data into actionable knowledge relies on technologies such as data-mining, machine learning, simulations, and context-awareness (Kessler & Löwe, 2012; Österlund & Löwe, 2018). We highlight two challenges on data driven technologies connected to the objectives of smarter systems.

The first challenge is to ensure *guarantees for properties* of systems that rely on data driven models throughout their lifetime. Data driven applications often outperform traditional ones in performance and accuracy on average. While this is good enough for many scenarios, some others need stricter guarantees even in the corner cases. For instance, a real-time control applications must guarantee a response before the deadline that is set by physical timing constraints; the control output signals must guarantee the stability of the controlled system for all observed input etc. This is difficult if the transfer function is complex, e.g., as a result of deep learning with several thousands of parameters. Verified AI (Seshia & Sadigh, 2016) aims to validate and guarantee system capabilities for all input, even in the corner cases. Moreover, when a continuous learning loop adapts the application based on new data, it is difficult to maintain and guarantee the capabilities of the application. Again, this applies to the functional correctness of the system and also to non-functional properties such as response time, security, and safety. The ability to ensure guarantees for functional and non-functional properties supports compliance of smarter systems with their requirements and maintain their resilience during operation, contributing to trustworthiness. The ability to maintain the guarantees over time makes smarter systems more sustainable.

The second challenge is to provide methods and techniques for *understanding data driven models* to support data driven applications and to transfer knowledge to other settings. It is related to the former but centred more around the human users and system engineers. Nowadays, it is textbook knowledge that the power of data-driven approaches is inversely proportional to their interpretability, cf. James, Witten, Hastie, and Tibshirani (2014). For instance, a linear regression model in one variable is less powerful and accurate, but easier to understand and to interpret, than a deep learning model in hundreds of variables with dozens of layers and thousands of parameters. Therefore, Explainable Artificial Intelligence (XAI) (Došilović, Brčić, & Hlupić, 2018; Gunning & Aha, 2019; Turek, 2016) aims at making even the accurate data-driven models understandable, interpretable, and trustable. More specifically, it "produces more explainable models, while maintaining a high level of learning performance (prediction accuracy), and enables human users to understand, appropriately trust, and effectively manage the emerging generation of artificially intelligent partners." This would also allow to generate repeatable actionable engineering knowledge (human intelligence) that can be transferred to new application cases and areas. Tackling this second challenge will be pivotal in coordinating the runtime activities of system adaptation with the human engineering activities for evolving smarter systems. This contributes to the sustainability of smarter systems.

### 3.4. Visual analytics

Visual analytics, which is defined as "the science of analytical reasoning facilitated by interactive visual interfaces" (Thomas & Cook, 2005), investigates ways to better comprehend large and complex data by combining the strengths of human and computational data processing. Visual analytics plays a crucial role in achieving a seamless integration of continuous adaptation and evolution that is required to mitigate uncertainty and ensuring stakeholder requirements in a trustworthy and sustainable manner across the lifetime of the system. The field of visualisation can be subdivided into two main sub-fields. On the one hand, scientific visualisation that focuses on visualising 3-dimensional data and temporal processes; here the spatial aspects of the data are crucial to correctly reflect the positions of the visualised real-world objects. On the other hand, information visualisation that focuses on abstract data (e.g., multidimensional or network data) and its visualisation through a visual mapping process. The resulting interactive visual representation should represent the abstract input data and support sense-making of the data, regardless of the many different data types, such as plain text, spatial, temporal and network data (Kerren & Schreiber, 2012). Besides visualisation, human visual perception and cognition are key areas of visual analytics that are naturally human-centred. They investigate the effects that user interfaces have on the analytics process' results, that is, to what degree the interface supports the user in successfully completing the analysis goal at hand. We highlight two key challenges of visual analytics in the following.

A first challenge is to *deal with specific characteristics of data of cyber-physical systems*, such as uncertainty in the data, the sheer size and complexity of real-world data in this context, and the open-endedness of these systems. These issues in combination with temporal aspects (e.g., introduced by streaming data) put high demands on visual analytics solutions but will be pivotal in the design and application of engineering processes for perpetual adaptation and evolution as required in smarter systems. Visualisation systems for analysing large-scale online social media text data are excellent showcases for addressing this specific challenge (Kucher, Paradis, Sahlgren, & Kerren, 2017). Tackling this first challenge contributes to configurable and adaptive life-cycle processes for smarter systems, enhancing the sustainability of smarter systems.

The second challenge is to ensure *interpretable and explainable machine learning models* with the help of interactive visualisation (Liu, Wang, Liu, & Zhu, 2017; Sacha et al., 2017) (cf. explainable AI above). Visual analytics tools for machine learning will contribute to increase the interpretability, explainability, and trust into such methods (Chatzimparmpas et al., 2020). An example is the t-viSNE approach for interactive assessment and interpretation of t-SNE projections (Chatzimparmpas, Martins, & Kerren, 2020). Tackling the second challenge will support smarter systems to communicate experiences in an understandable manner to engineers, users, and other stakeholders, where visual analytics can play a key role in achieving this goal. Tackling this second challenge will be pivotal in the seamless integration of evolution and adaptation activities, contributing both to the trustworthiness and sustainability of smarter systems.

### 3.5. Synergies between core areas

As explained, we have selected cyber-physical systems, self-adaptation, data driven technology, and visual analytics as core areas for future research on smarter systems. The choice for these areas is motivated by their mapping to key characteristics and demands of smarter systems. Yet, we acknowledge that this choice implies a particular viewpoint on the research challenges for smarter systems. Our aim is not to be exclusive; other viewpoints can be defined starting from different angles that would provide complementary challenges. In this section, we outline principle connections between cyber-physical systems, self-adaptation, data driven technology, and visual analytics that lead to synergies for creating smarter systems. Our aim is not to be exhaustive, but rather to highlight examples that illustrate (some of) these synergies.

**CPS and data driven technologies.** Model-based simulations of cyber-physical systems can create huge amounts of labelled data at low costs. Such data can be pivotal for employing more advanced learning technologies, such as deep learning. More precisely, simulations can be conducted under controlled conditions (the ground truth labelling) to generate observations. Deep learning can then map these and similar real-world observations to the similar root cause conditions. For example, in the Smart Grid scenario, faults can be injected at different segments of the power lines and the corresponding simulated time series observations at all sensors can be captured. This generates a massive amount of data sufficient for training a deep learning model for the inverse mapping, i.e., mapping sensor observations to fault locations.[1]

**Visual analytics for data driven technologies.** A challenge with machine learning is that approaches with higher predictive power, such as deep learning, provide less human-understandable explanations for the underlying phenomena and vice versa. Also, finding the right setup of predictors and hyper-parameters needs human insights. Visual analytics allows creating human insights in complex (technical) phenomena that are connected with huge data sets. This is a way to improve the understandability of and, hence, establishing trust for stakeholders in (deep) learning models and reducing the engineering effort for developing and training data driven models.

**Self-adaptation for data driven technologies.** Hyper-parameter optimisation and feature engineering in machine learning are mainly human efforts today. They are partially supported by automated approaches, such as systematic optimisation (and even a self-application of machine learning). There exists, however, also an opportunity to generally understand these processes as self-adaptation, where machine learning is the managed system for which hyper-parameters and features need to be selected and self-adapted, based on observations such as amount and kinds of data available, current loss, etc.

[1] For an example, see the Linnaeus University Centre for Data Intensive Sciences and Applications: https://lnu.se/forskning/sok-forskning/linnaeus-university-centre-for-data-intensive-sciences-and-applications/saddprojekt/ground-fault-location/.

**Data driven technology for self-adaptation.** Self-adaptation maps monitored observations to change actions that, in turn, can later be observed being successful to some degree. Central to this perpetual loop is the creation of knowledge that is used by the self-adaptation feedback loop to reason. The natural way of automatically creating knowledge from observations is machine learning, i.e., presumably successful actions as a reaction on yet unseen observations can be trained using supervised and feedback learning. Delegating knowledge acquisition and pre-processing of large amounts of data for self-adaptation to data-driven technologies will be key in data-intensive domains as we envision in smarter systems.

**Data driven discovery for CPS.** The integration of cyber and physical components complicates the design of systems that demand high utility – i.e., performance, safety, security, sustainability (i.e., near-zero power consumption), scalability, usability. Data driven technology can play a key role in inferring cyber-physical system models from data. Models (usually expressed in terms of laws and equations) facilitate the understanding of complex phenomena and allow for systems analysis and predictions. However, explicit modelling of cyber-physical systems may suffer from their intrinsic complexity arising from the combination and interaction of cyber and physical components. Data-driven technologies can facilitate the data-to-discovery process of implicit cyber-physical system Models. Specifically, employing Data-driven technologies allow for learning and inferring implicit models from the observation of raw data. These technologies have the advantage of (i) testing correlations between different variables and observations, and (ii) learning unforeseen patterns in system behaviours.

**Data driven analysis for CPS.** Data driven technologies also provide the means for inferring quality properties from cyber-physical system models and runtime data. Analysing, assessing, and preserving quality goals is a key concern of cyber-physical systems. To that end, cyber-physical system models (implicit or explicit) can be fed with data monitored from the cyber-physical system constituents and then used for assessing the level of actual quality, as well as for forecasting the future one. When deviations to the expected behaviour are detected/forecasted, the system may be adapted in order to preserve the quality goals – e.g., performance, safety, security, sustainability. Hence, analysis must be performed on-line to continuously assess the current behaviour in a timely and efficient manner. It is well known that the computational complexity of model-based analysis techniques is one of the key challenges in system verification. To that end, data driven technologies, such as for instance search-based techniques, reinforcement learning, and neural networks, enable online model-based prediction and estimation.

**Process mining for CPS with data driven technologies and Self-adaptation.** Process Mining couples data science with process science, offering a valuable alternative to detect and predict faults or anomalies in cyber-physical systems. Starting from event data, process mining has three intended uses: (1) process discovery, where new process models can be discovered; (2) conformance check, where one or more existing process models can be replayed to verify the conformance of the behaviour of the application with the considered process models; (3) enhancement, where data can be replayed on the models to perform performance and dependability analyses. When mining (discovering) a process model, different perspectives may be captured. One is the control-flow perspective, where the ordering of activities is found and recurrent patterns are captured by a model. Applied to cyber-physical system resilience, process mining can support anomaly detection from event data coming from the edge, based on process discovery from available data. After a model is retrieved, additional data is collected and each trace of related activities in the data is replayed on the model. If differences are detected a fitness parameter is determined to compare it with a specified threshold. Based on an assessment the cyber physical system may then be classified as misbehaving.

**Combining the core areas for auto-evolving CPS.** Maintaining and evolving cyber-physical systems is a challenging task. While tool support is increasing at high speed to manage complexity and reducing turn-around time, in essence maintenance and evolution remain a human-dependent activity. Combining self-adaptation with data driven technologies and visual analytics has the potential to realise a breakthrough in the way maintenance and evolution are realised today. This will require lifting self-adaptation up to the level of self-learning (in particular evolutionary learning) where on the one hand data driven technologies are exploited to manage huge amounts of data, and visual analytics on the other hand enable taking the human in the loop. Key aspects will be the provision and management of computing resources and service data that need to be integrated automatically over time as well as the management of historical data.

## 4. Research Agenda

We bring now together the key research objectives for capabilities of smarter systems and their engineering processes with the challenges of key areas to achieve these objectives. This allows us to devise a multi-year research agenda for trustworthy and sustainable smarter systems. Realising this research agenda will require a concerted effort of multiple research teams active in one or more of the four core areas. This section provides an outline of this research agenda for the next decade that we centre around three themes: assurances for unknowns, self-explainability, and smarter ecosystems, see also Figure 1.

It is important to note that the proposed research agenda does not aim to propose a blue-sky vision based on radically novel ideas. Instead, the research agenda targets complex challenges grounded in scientific insights and emerging principles obtained from the state-of-the-art.

Planning research over a period of a decade is obviously a difficult and risky task. After all, research objectives only provide the drivers for research endeavours and intermediate results will inevitably further shape the research scope and its direction. We split up each theme in two parts: the first part covering more concrete lines of research for the first five years, while the second covers more speculative research for the next five years. Figure 2 provides a schematic overview of the research agenda.

### 4.1. Theme 1: Assurances for unknowns

The first theme is centred on a line of integrated research that focuses on the first key objective of assuring system capabilities in uncertain conditions. Theme 1 concerns the enhancements of smarter systems in their functional capacity and utility in the face of a non-fully deterministic future of smarter systems.

*Part 1: Assurances for Decentralised Cyber-Physical-Systems*. The first part proposes the study of assurances for the behaviour of decentralised cyber-physical systems that operate under uncertainty. Decentralisation refers to the integration of multiple decision-making entities. Decentralisation is getting increasingly important; it may be required for quality purposes (e.g., to achieve a scalable solution or avoid a single point of failure), or it may be implied by the nature of the problem domain (e.g., systems that cross ownership domains). The aim is to investigate stakeholder concerns, goal models, assurance structures (assurance cases (Calinescu et al., 2018)), and supporting data analytic techniques. Central to the study are mechanisms for coordinating online and offline reasoning and decision-making across multiple entities with trade-offs for conflicting goals, the required verification with consensus, and coordinated enactment of system adaptation and evolution. An important aspect will be the management of emergent behaviour. This may call for mechanisms to detect deviations of regular system behaviour, such as anomaly detection (Bhuyan, Bhattacharyya, & Kalita, 2014). There is a need to investigate the role of visual analytics in verification activities and the communication of verification results to stakeholders in a proper way, for instance in the context of a dynamic certification process. A key deliverable will be an approach for dynamic assurance cases.

*Part 2: Assurances for Smarter Cyber-Physical-Systems with Unknowns*. Followup research proposes to

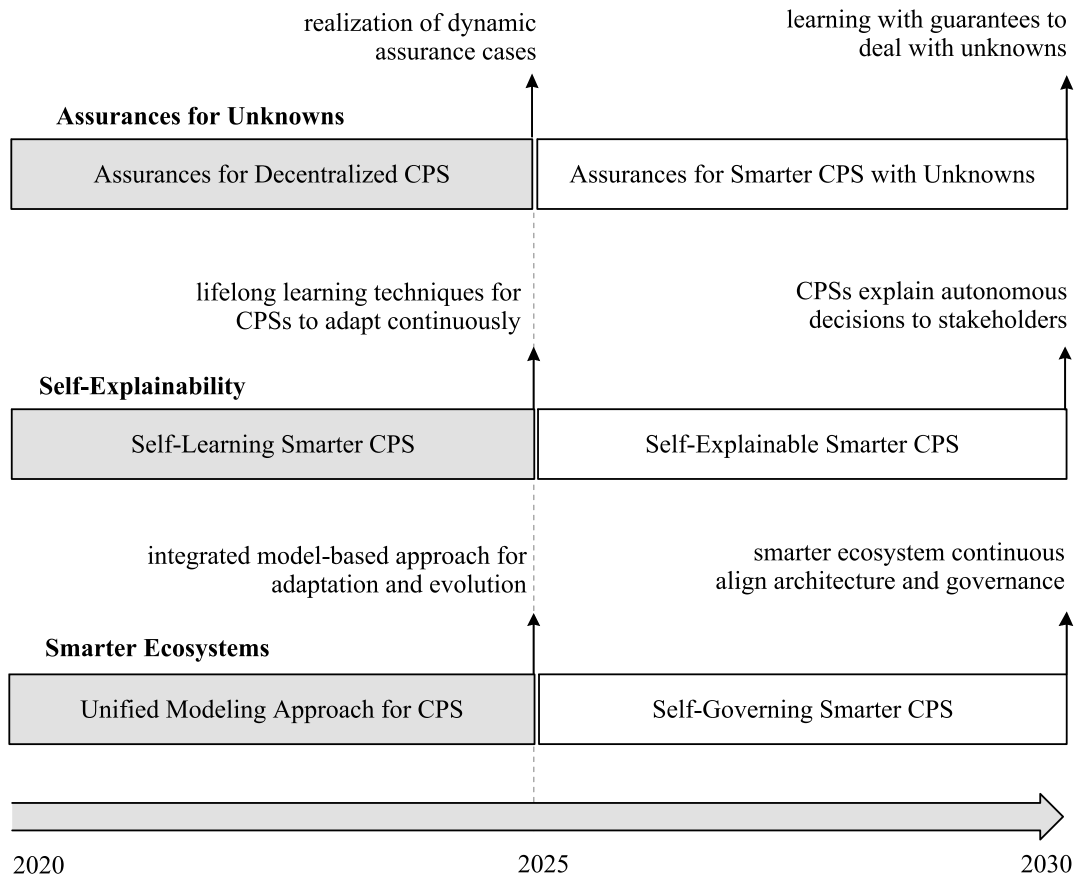


**Fig. 2. Planning of Research. Gray Boxes Represent a First Iteration of Research with Examples of Concrete Milestones; White Boxes Represent Follow up Research.**

investigate how to establish assurances for smarter cyber-physical systems that are subject of unknowns, i.e., conditions the system faces during operation that were not completely anticipated. Manually reacting to such novel situations would require tremendous efforts and may be too slow in critical situations. Dealing with these intrinsic challenges requires a cyber-physical system to preserve knowledge from the past and utilise this knowledge efficiently when performing tasks in the future. This calls for the investigation of unsupervised learning techniques that enable automated, or where needed human-supported, discovery of novelty. Interesting approaches in this direction could be deep learning (Goodfellow, Bengio, & Courville, 2016) and so called subspace clustering (Chen, Lv, & Yi, 2021; Vidal, 2011). The second part studies different online learning techniques with a particular focus on the boundaries of guarantees that such techniques can offer. A key deliverable will be learning techniques with guarantees to deal with unknowns.

## 4.2. Theme 2: Self-explainability

The second theme is centred on a line of integrated research that focuses on both key objectives: assuring system capabilities in uncertain conditions, and engineering processes for perpetual adaptation and evolution. Theme 1 concerns the ability of smarter systems to communicate and interact with humans.

*Part 1: Self-Learning Cyber-Physical Systems.* The types of uncertainties that next generation cyber-physical systems will face, as well as the amount of data they will need to process, requires such systems to learn over time. Leveraging on recent progress in the field of machine learning, this part proposes the study of continuous lifelong learning (Chen & Liu, 2020)—referring to the cyber-physical systems' lifetime—for next generation cyber-physical systems. In particular, this part studies lifelong meta-learning that allows cyber-physical systems to learn a learning approach from many related tasks. To this end, relevant cyber-physical systems or their constituents need to be enhanced with a meta-learning system that offers: a) facilities to store and manage knowledge, b) a scheme to represent relevant knowledge, c) a meta-learner that initiates and evolves a learner from experiences of executed tasks, and d) a learner that exploits the learned knowledge to deal with new learning tasks and emerging situations that the component or system encounters. A key deliverable will be a lifelong learning approach for smarter cyber-physical systems.

*Part 2: Self-Explainable Cyber-Physical-Systems.* The complexity of future cyber-physical systems, induced by the need for integrated adaptation and evolution, with human involvement makes it practically impossible for stakeholders to understand the system structure and behaviour at all times. Followup research of this theme proposes to investigate techniques and tools that add a self-explainable capability to a cyber-physical system. The capability provides for external queries regarding, among other, system structure and behaviour, decisions and rationale, and system and environment state. A future cyber-physical system makes decisions based on data-analysis and AI techniques. A self-explainable system can communicate such a process, a rationale for decisions and their outcome to external parties, which is a prerequisite to establishing trust. Concrete challenges include: comprehensibility of explanations, the presentation of explanations, human-machine interactivity and conversations, and a-posterior explanations (Blumreiter et al., 2019). A key deliverable will be an approach that equips smarter cyber-physical systems with self-explainability capabilities.

### 4.3. Theme 3: Smarter ecosystems

This theme is centred around a line of integrated research that focuses on the key objective of engineering processes for perpetual adaptation and evolution.

*Part 1: Unified Modelling Approach for Smarter Cyber-Physical Systems.* The first part proposes to investigate models, techniques, and tools for smarter engineering processes. Models and model transformations will drive a smarter system and its smarter engineering process. The projects investigates design and runtime models that are involved in the specification, adaptation, and evolution of smarter cyber-physical systems. Smarter cyber-physical systems will typically be built by multi-disciplinary teams that use heterogeneous models that vary in formalism's, concepts, and levels of abstraction. This calls for a fluid modelling approach that combines different specification and verification approaches (Ruchkin, De Niz, Chaki, & Garlan, 2014). Adopting a single, all-encompassing homogeneous modelling language denies the fact that domain-specific formalism's are better at verifying properties for their domain, and that there is usually a well-established body of knowledge and expertise built up around these formalism's. Hence, domain specific languages and developer tools will be essential for creating, integrating and maintaining models. Central will be proper definition of modelling abstractions with first-class support for multi-model integration, properties to express relationships between the abstractions of models, and means to execute domain-specific analyses on models across the lifetime of the system. A key deliverable will be an integrated model-based framework for adaptation and evolution.

*Part 2: Self-Governing Smarter Cyber-Physical Systems.* Future cyber-physical systems will form ecosystems, where system owners and third-parties share responsibilities. The long term objective of this part is to support the seamless integration of the engineering and operation of smarter ecosystems leveraging on a model-driven approach. This effort proposes the study of novel languages, models, and transformations, but also supporting infrastructure and tools that allow two-way interaction between the ecosystem and its stakeholders, including developers, operators, and users. An ecosystem requires intentional, not ad-hoc, management of ecosystem partners (Bosch, 2016). Smarter cyber-physical ecosystems will expose a high level of autonomy requiring reflective capabilities where the system collects data about its utility and adjusts according to its goals. Because data will be produced by different parts of the system semantically underpinned data-fusion techniques will be required. A key deliverable will be a smarter ecosystem that continuously aligns its architecture and governance to the concerns of a variety of stakeholders.

## 5. Conclusions

Ensuring the required trustworthiness and sustainability of systems that blend cyber, physical, and social elements will be vital for our society. Yet, engineering such systems is complex. Primary reasons for this complexity are the uncertainty and continuously change systems face, the presence of large volumes of data that needs to be processed, and the role of humans as inherent parts of the systems. To tackle this complexity, we argued that both systems and the way we engineer them must become smarter, meaning that both systems and the processes to engineer them adapt and evolve through a perpetual and enduring process that continuously improves their capabilities to deal with the uncertainties and change they face across their lifetime. We proposed key objectives for engineering smarter systems and highlighted core areas that are expected to be pivotal in achieving these goals together with their respective challenges. From this, we proposed an ambitious research agenda for smarter systems for the next decade. Key research targets that are centred around three themes: assurances for unknowns, self-explainability, and smarter ecosystems are:

- An approach for dynamic assurance cases to provide assurances for cyber-physical-systems in decentralised settings.
- Online learning techniques with guarantees to provide assurances for smarter cyber-physical systems that have to deal with unknowns.
- A lifelong learning approach that enables smarter cyber-physical systems to deal with new tasks and novel emerging situations.
- Self-explainability capabilities for smarter cyber-physical systems enacting seamless integration of human operators.
- An integrated multi-model based framework for adaptation an evolution that spans the full life-cycle of smarter systems.
- The foundations for a smarter ecosystem that continuously aligns its architecture and governance to the concerns of a variety of stakeholders.

Realising this research agenda will require a multi-year concerted effort of research teams active in the different core areas of smarter systems. We hope that this paper will offer a source of inspiration for those who want to study and develop novel solutions for trustworthy and sustainable computing systems to the good of our society.

## References

Andersson, J., Baresi, L., Bencomo, N., de Lemos, R., Gorla, A., Inverardi, P., & Vogel, T. (2013). Software engineering processes for self-adaptive systems. In R. de Lemos, H. Giese, H. A. Müller, & M. Shaw (Eds.), *Software Engineering for Self-Adaptive Systems II* (pp. 51–75). Springer. doi: 10.1007/978-3-642-35813-5˙3

Baheti, R., & Gill, H. (2019). Cyber-physical systems. In *IEEE International Conference on Mechatronics* (Vol. 1, p. 430-432). doi: 10.1109/ICMECH.2019.8722929

Banerjee, A., et al. (2012). Ensuring safety, security, and sustainability of mission-critical cyber–physical systems. *Proceedings of the IEEE*, *100*(1), 283-299. doi: 10.1109/JPROC.2011.2165689

Baresi, L., & Ghezzi, C. (2010). The disappearing boundary between development-time and run-time. In *Future of Software Engineering Research* (p. 17–22). ACM. doi: 10.1145/1882362.1882367

Becker, C., Chitchyan, R., Duboc, L., Easterbrook, S., Penzenstadler, B., Seyff, N., & Venters, C. C. (2015). Sustainability design and software: The karlskrona manifesto. In *International Conference on Software Engineering* (Vol. 2, p. 467-476). doi: 10.1109/ICSE.2015.179

Bennaceur, A., Ghezzi, C., Tei, K., Kehrer, T., Weyns, D., Calinescu, R., ... Zhao, H. (2019). Modelling and analysing resilient cyber-physical systems. In *14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems* (pp. 70–76). ACM. doi: 10.1109/SEAMS.2019.00018

Bhuyan, M. H., Bhattacharyya, D. K., & Kalita, J. K. (2014). Network anomaly detection: Methods, systems and tools. *IEEE Communications Surveys Tutorials*, *16*(1), 303-336. doi: 10.1109/SURV.2013.052213.00046

Blair, G., Bencomo, N., & France, R. B. (2009). Models@ run.time. *Computer*, *42*(10), 22–27. doi: 10.1109/MC.2009.326

Blumreiter, M., Greenyer†, J., Javier, F., Garcia, C., Klos, V., Schwammbergerk, M., ... Wortmann, A. (2019). Towards self-explainable cyber-physical systems. *arXiv:1908.04698v1*, 1-6.

Bosch, J. (2016). Speed, data, and ecosystems: The future of software engineering. *IEEE Software*, *33*(1), 82-88. doi: 10.1109/MS.2016.14

Bures, T., et al. (2018). Software engineering for smart cyber-physical systems: Models, system-environment boundary, and social aspects. *ACM SIGSOFT Software Engineering Notes*, *43*(4), 42–44. doi: 10.1145/3282517.3302401

Bures, T., Weyns, D., Berger, C., Biffl, S., Daun, M., Gabor, T., ... Pronios, N. (2015). Software engineering for smart cyber-physical systems – towards a research agenda. *ACM SIGSOFT Software Engineering Notes*, *40*(6), 28–32. doi: 10.1145/2830719.2830736

Calinescu, R., Mirandola, R., Perez-Palacin, D., & Weyns, D. (2020). Understanding uncertainty in self-adaptive systems. In *2020 ieee international conference on autonomic computing and self-organizing systems (acsos)* (p. 242-251). doi: 10.1109/ACSOS49614.2020.00047

Calinescu, R., Weyns, D., Gerasimou, S., Iftikhar, M., Habli, I., & Kelly, T. (2018). Engineering trustworthy self-adaptive software with dynamic assurance cases. *IEEE Transactions on Software Engineering*, *44*(11).

Caporuscio, M., Flammini, F., Khakpour, N., Singh, P., & Thornadtsson, J. (2020). https://doi.org/10.1016/j.future.2019.09.004.smart-troubleshooting connected devices: Concept, challenges and opportunities. *Future Generation Computer Systems*, *111*, 681-697. doi: 10.1016/j.future.2019.09.004.

Chatzimparmpas, A., Martins, R. M., Jusufi, I., Kucher, K., Rossi, F., & Kerren, A. (2020). The state of the art in enhancing trust in machine learning models with the use of visualizations. *Computer Graphics Forum*, *39*(3), 713–756. doi: 10.1111/cgf.14034

Chatzimparmpas, A., Martins, R. M., & Kerren, A. (2020). t-viSNE: Interactive assessment and interpretation of t-SNE projections. *IEEE Transactions on Visualization and Computer Graphics*, *26*(8), 2696–2714. doi: 10.1109/TVCG.2020.2986996

Chen, D., Lv, J., & Yi, Z. (2021). Unsupervised multi-manifold clustering by learning deep representation. *Crowdsourcing, Deep Learning, and Artificial Intelligence Agents*.

Chen, Z., & Liu, B. (2020). *Lifelong machine learning*. Morgan and Cleypool Publishers.

Cheng, B. H. C., Ramirez, A., & McKinley, P. K. (2013). Harnessing evolutionary computation to enable dynamically adaptive systems to manage uncertainty. In *International Workshop on Combining*

*Modelling and Search-Based Software Engineering* (p. 1-6). doi: 10.1109/CMSBSE.2013.6604427

Ciccozzi, F., Tichy, M., Vangheluwe, H., & Weyns, D. (2019). Blended modelling - what, why and how. In *ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems Companion* (p. 425-430). doi: 10.1109/MODELS-C.2019.00068

CPSWeek. (2021). Cyber-physical systems.

D'Angelo, M., Napolitano, A., & Caporuscio, M. (2018). Cyphef: A model-driven engineering framework for self-adaptive cyber-physical systems. In *40th International Conference on Software Engineering: Companion* (p. 100-104). ACM. doi: 10.1145/3183440.3183483

Došilović, F. K., Brčić, M., & Hlupić, N. (2018). Explainable artificial intelligence: A survey. In *41st International Convention on Information and Communication Technology, Electronics and Microelectronics* (p. 210-215). doi: 10.23919/MIPRO.2018.8400040

Esfahani, N., & Malek, S. (2013). Uncertainty in self-adaptive software systems. In R. de Lemos, H. Giese, H. A. Müller, & M. Shaw (Eds.), *Software engineering for self-adaptive systems ii: International seminar, dagstuhl castle, germany, october 24-29, 2010 revised selected and invited papers* (pp. 214–238). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-642-35813-5˙9

EU. (2021a). Digital europe for a more competitive, autonomous and sustainable europe.

EU. (2021b). H2020 Program Cyber Physical Systems.

Fitz, T., Theiler, M., & Smarsly, K. (2019). A metamodel for cyber-physical systems. *Advanced Engineering Informatics*, *41*, 100930. doi: 10.1016/j.aei.2019.100930

Flammini, F., Marrone, S., Nardone, R., Caporuscio, M., & D'Angelo, M. (2020). Safety integrity through self-adaptation for multi-sensor event detection: Methodology and case-study. *Future Generation Computer Systems*, *112*, 965-981. doi: https://doi.org/10.1016/j.future.2020.06.036

Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management*, *35*(2), 137-144. doi: https://doi.org/10.1016/j.ijinfomgt.2014.10.007

Garlan, D., Cheng, S.-W., Huang, A.-C., Schmerl, B., & Steenkiste, P. (2004). Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, *37*(10), 46-54. doi: 10.1109/MC.2004.175

Gheibi, O., Weyns, D., & Quin, F. (2021). Applying machine learning in self-adaptive systems: A systematic literature review. *ACM Transactions on Autonomous and Adaptive Systems*.

Gol Mohammadi, N. (2019). *Trustworthy cyber-physical systems - A systematic framework towards design and evaluation of trust and trustworthiness*. Springer. ISBN 978-3-658-27488-7.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. ISBN 978-0262035613.

Gunes, V., Peter, S., Peter, S., Givargis, T., Vahid, F., & Vahid, F. (2014). A survey on concepts, applications, and challenges in cyber-physical systems. *KSII Transactions on Internet and Information Systems*, *8*(12), 4242-4268. doi: 10.3837/tiis.2014.12.001

Gunning, D., & Aha, D. (2019). Darpa's explainable artificial intelligence (XAI) program. *AI Magazine*, *40*(2), 44-58. doi: 10.1609/aimag.v40i2.2850

Gupta, S. K., et al. (2011). Research directions in energy-sustainable cyber–physical systems. *Sustainable Computing: Informatics and Systems*, *1*(1), 57-74. doi: 10.1016/j.suscom.2010.10.003

IBM. (2003). An architectural blueprint for autonomic computing.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2014). *An introduction to statistical learning: With applications in R*. Springer. ISBN 1461471370.

Jazdi, N. (2014). Cyber physical systems in the context of industry 4.0. In *IEEE International Conference on Automation, Quality and Testing, Robotics* (p. 1-4). doi: 10.1109/AQTR.2014.6857843

Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, *36*(1), 41-50. doi: 10.1109/MC.2003.1160055

Kephart, J. O., Dibia, V. C., Ellis, J. B., Srivastava, B., Talamadupula, K., & Dholakia, M. (2019). An embodied cognitive assistant for visualizing and analyzing exoplanet data. *IEEE Internet Comput.*, *23*(2), 31–39. doi: 10.1109/MIC.2019.2906528

Kerren, A., & Schreiber, F. (2012). Toward the role of interaction in visual analytics. In *Winter Simulation Conference* (pp. 420:1–420:13). IEEE Computer Society Press. doi: 10.1109/WSC.2012.6465208

Kessler, C., & Löwe, W. (2012). Optimized composition of performance-aware parallel components. *Concurrency and Computation: Practice and Experience*, *24*(5), 481–498. doi: 10.1002/cpe.1844

Koutsoukos, X., Karsai, G., Laszka, A., Neema, H., Potteiger, B., Volgyesi, P., … Sztipanovits, J. (2018). Sure: A modeling and simulation integration platform for evaluation of secure and resilient cyber–physical systems. *Proceedings of the IEEE*, *106*(1), 93-112. doi: 10.1109/JPROC.2017.2731741

Kramer, J., & Magee, J. (2007). Self-managed systems: An architectural challenge. In *Future of Software Engineering* (p. 259-268). IEEE. doi: 10.1109/FOSE.2007.19

Kucher, K., Paradis, C., Sahlgren, M., & Kerren, A. (2017). Active learning and visual analytics for stance classification with alva. *ACM Transactions on Interactive Intelligent Systems*, *7*(3), 31. doi: 10.1145/3132169

Lee, E. A. (2008). Cyber Physical Systems: Design challenges. In *IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing* (p. 363-369). doi: 10.1109/ISORC.2008.25

Lee, J., Bagheri, B., & Kao, H.-A. (2015). A cyber-physical systems architecture for industry 4.0-based manufacturing systems. *Manufacturing Letters*, *3*, 18-23. doi: 10.1016/j.mfglet.2014.12.001

Lin, Y., Wang, P., & Ma, M. (2017). Intelligent transportation system: Concept, challenge and opportunity. In *International Conference on Big Data Security on Cloud, International Conference on High Performance and Smart Computing, and International Conference on Intelligent Data and Security* (p. 167-172). doi: 10.1109/BigDataSecurity.2017.50

Liu, S., Wang, X., Liu, M., & Zhu, J. (2017). Towards better analysis of machine learning models: A visual analytics perspective. *Visual Informatics*, *1*(1), 48–56. doi: 10.1016/j.visinf.2017.01.006

Liu, Z., Yang, D., Wen, D., Zhang, W., & Mao, W. (2011). Cyber-physical-social systems for command and control. *IEEE Intelligent Systems*, *26*(4), 92-96. doi: 10.1109/MIS.2011.69

Mahdavi-Hezavehi, S., Avgeriou, P., & Weyns, D. (2017). Chapter 3 - a classification framework of uncertainty in architecture-based self-adaptive systems with multiple quality requirements. In I. Mistrik, N. Ali, R. Kazman, J. Grundy, & B. Schmerl (Eds.), *Managing trade-offs in adaptable software architectures* (p. 45-77). Morgan Kaufmann. doi: https://doi.org/10.1016/B978-0-12-802855-1.00003-4

NSF. (2021). National Science Foundation 21-551, Cyber Physical Systems..

Oreizy, P., Gorlick, M. M., Taylor, R. N., Heimhigner, D., Johnson, G., Medvidovic, N., … Wolf, A. L. (1999). An architecture-based approach to self-adaptive software. *IEEE Intelligent Systems*, *14*(3), 54-62. doi: 10.1109/5254.769885

Österlund, E., & Löwe, W. (2018). Self-adaptive concurrent components. *Automated Software Engineering*, *25*(1), 47–99. doi: 10.1007/s10515-017-0219-0

Pagliari, L., Mirandola, R., & Trubiani, C. (2020). Engineering cyber-physical systems through performance-based modelling and analysis: A case study experience report. *Journal of Software: Evolution and Process*, *32*(1), e2179. doi: 10.1002/smr.2179

Ratasich, D., Khalid, F., Geissler, F., Grosu, R., Shafique, M., & Bartocci, E. (2019). A roadmap toward the resilient internet of things for cyber-physical systems. *IEEE Access*, *7*, 13260-13283. doi: 10.1109/ACCESS.2019.2891969

Ruchkin, I., De Niz, D., Chaki, S., & Garlan, D. (2014). Contract-based integration of cyber-physical analyses. In *International Conference on Embedded Software* (p. 1-10). doi: 10.1145/2656045.2656052

Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Peltonen, J., Weiskopf, D., … Keim, D. A. (2017). What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing*, *268*, 164–175. doi: https://doi.org/10.1016/j.neucom.2017.01.105

Schmidt, M., & Åhlund, C. (2018). Smart buildings as cyber-physical systems: Data-driven predictive control strategies for energy efficiency. *Renewable and Sustainable Energy Reviews*, *90*, 742-756. doi: https://doi.org/10.1016/j.rser.2018.04.013

Selic, B. (2020). Controlling the controllers: What software people can learn from control theory. *IEEE Softw.*, *37*(6), 99–103. doi: 10.1109/MS.2020.3006970

Seshia, S. A., & Sadigh, D. (2016). Towards verified artificial intelligence. *CoRR*, *abs/1606.08514*.

Stankovic, J. A. (2016). Research directions for cyber physical systems in wireless and mobile healthcare. *ACM Transactions on Cyber-Physical Systems*, *1*(1), 1-12. doi: 10.1145/2899006

Sztipanovits, J., Koutsoukos, X., Karsai, G., Kottenstette, N., Antsaklis, P., Gupta, V., ... Wang, S. (2012). Toward a science of cyber–physical system integration. *Proceedings of the IEEE*, *100*(1), 29-44. doi: 10.1109/JPROC.2011.2161529

Tao, F., Qi, Q., Wang, L., & Nee, A. (2019). Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison. *Engineering*, *5*(4), 653-661. doi: 10.1016/j.eng.2019.01.014

Tavčar, J., & Horváth, I. (2019). A review of the principles of designing smart cyber-physical systems for run-time adaptation: Learned lessons and open issues. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *49*(1), 145-158. doi: 10.1109/TSMC.2018.2814539

Thomas, J. J., & Cook, K. A. (2005). *Illuminating the path: The research and development agenda for visual analytics*. IEEE CS Press.

Tokody, D., Papp, J., Iantovics, L. B., & Flammini, F. (2019). Complex, resilient and smart systems. In *Resilience of cyber-physical systems* (p. 3-24). Springer. doi: 10.1007/978-3-319-95597-1˙1

Tuballa, M. L., & Abundo, M. L. (2016). A review of the development of smart grid technologies. *Renewable and Sustainable Energy Reviews*, *59*, 710-725. doi: https://doi.org/10.1016/j.rser.2016.01.011

Turek, M. (2016). *Explainable artificial intelligence (XAI).* DARPA.

Van Der Donckt, J., Weyns, D., Quin, F., Van Der Donckt, J., & Michiels, S. (2020). Applying deep learning to reduce large adaptation spaces of self-adaptive systems with multiple types of goals. In *Software Engineering for Adaptive and Self-Managing Systems* (p. 20–30). IEEE. doi: 10.1145/3387939.3391605

Vidal, R. (2011). Subspace clustering. *IEEE Signal Processing Magazine*, *28*(2), 52-68. doi: 10.1109/MSP.2010.939739

Weyns, D. (2019). Software engineering of self-adaptive systems. In S. Cha, R. N. Taylor, & K. C. Kang (Eds.), *Handbook of software engineering.* (pp. 399–443). doi: 10.1007/978-3-030-00262-6

Weyns, D. (2020). *Introduction to self-adaptive systems: A contemporary software engineering perspective*. Wiley, IEEE Computer Society Press. ISBN 978-1-119-57494-1.

Weyns, D., Caporuscio, M., Vogel, B., & Kurti, A. (2015). Design for sustainability = runtime adaptation U evolution. In *European Conference on Software Architecture Workshops* (p. 1-7). doi: 10.1145/2797433.2797497

Weyns, D., & Iftikhar, M. U. (2019). ActivFORMS: A Model-Based Approach to Engineer Self-Adaptive Systems. In *CoRR, abs/1908.11179.*

Yu, X., & Xue, Y. (2016). Smart grids: A cyber–physical systems perspective. *Proceedings of the IEEE*, *104*(5), 1058-1070. doi: 10.1109/JPROC.2015.2503119

Zeng, J., Yang, L. T., Lin, M., Ning, H., & Ma, J. (2020). A survey: Cyber-physical-social systems and their system-level design methodology. *Future Generation Computer Systems*, *105*, 1028-1042. doi: https://doi.org/10.1016/j.future.2016.06.034

Zhou, K., Liu, T., & Zhou, L. (2015). Industry 4.0: Towards future industrial opportunities and challenges. In *2015 12th international conference on fuzzy systems and knowledge discovery (fskd)* (p. 2147-2152). doi: 10.1109/FSKD.2015.7382284

## Author Biographies

**Danny Weyns** is a professor at the Katholieke Universiteit Leuven, Belgium and Linnaeus University, Sweden. He obtained a PhD from the Katholieke Universiteit Leuven for research on software architectures for multi-agent systems. His current research interests are in software engineering of self-adaptive systems.

**Jesper Anderson** is Associate Professor at Linnaeus University, Sweden. He received his PhD from Linköping University for research on dynamic software architectures. His current research interests are in software variability, systematic software reuse, and engineering processes for self-adaptive systems.

**Mauro Caporuscio** is Associate Professor at Linnaeus University, Sweden. He received his PhD in Computer Science from the University of L'Aquila, Italy. His research interests mainly focus on software engineering and self-adaptive systems, with particular emphasis on decentralisation and resiliency.

**Francesco Flammini** is a professor at Mälardalen University and Linnaeus University, Sweden. He holds a PhD from the University of Naples Federico II, Italy for research on model-based evaluation of critical control systems. His current research interests are in resilient CPS and trustworthy autonomy.

**Andreas Kerren** is a professor of information visualization at Linköping University (LiU) and Linnaeus University (Lnu), Sweden. He is head of the ISOVIS research group at Lnu and leader of the iVis group at LiU. His research interests mainly focus on visual network and text analytics, as well as explainable AI.

**Welf Löwe** is a professor in computer science at Linnaeus University (Lnu), Sweden. He is also co-founder/CSO of Softwerk AB and head of the research excellence centre on Data-Intensive Sciences and Application at Lnu. His research focuses on traditional and data-driven software construction technologies.