Requirements for an AI-enabled Industry 4.0 Platform - Integrating Industrial and Scientific Views

Eichelberger Holger
Software Systems Engineering
University of Hildesheim
Hildesheim, Germany
e-mail: eichelberger@sse.uni-hildesheim.de

Stichweh Heiko
Head of Innovation
Lenze SE
Hamelin, Germany
e:mail: heiko.stichweh@lenze.com

Sauer Christian
Software Systems Engineering
University of Hildesheim
Hildesheim, Germany
e-mail: sauer@sse.uni-hildesheim.de

***Abstract*—Intelligent manufacturing is one goal of smart industry/Industry 4.0 that could be achieved through Artificial Intelligence (AI). Flexibly combining AI methods and platform capabilities, such as dynamic offloading of code close to production machines, security or interoperability mechanisms are major demands in this context. However, recent Industry 4.0 software platforms fall short in various of these demands, in particular in upcoming ecosystem scenarios, e.g., when data or services shall be shared across platforms or companies without vendor lock-ins. The aim of the funded Intelligent Industrial Production (IIP) IIP-Ecosphere project is to research concepts and solutions for 'easy-to-use' AI in Industry 4.0 and to demonstrate the results in a prototypical software platform. Core questions are which demands shall drive the development of such a platform and how a feasible set of requirements can be determined that balances scientific and industrial interests. In this paper, we discuss our approach on eliciting requirements in this context for two interlinked requirements perspectives, a usage and a functional view. In summary, we collected 67 usage view activities / scenarios and 141 top-level requirements with 179 detailing sub-requirements. About 35% of the requirements have so far been realized in a prototype and some of the identified concepts are currently being taken up by a standardization initiative for edge devices in Industry 4.0.**

***Keywords*—Industry 4.0 platforms; intelligent production; AI; requirements; edge; adaptation; asset administration shell.**

I. INTRODUCTION

The digitization of industry increases the performance of technical systems and their processes, but also their complexity. Intelligent manufacturing (smart industry, Industry 4.0) can be realized through application of Artificial Intelligence (AI) in the production context. This is perceived as an enabler for an increase of productivity of up to 50% [1]. However, currently more than 75% of AI applications are ultimately not deployed [2], e.g., as they are not considered to be production ready or as they are not easily applicable by domain users.

One further trend in Industry 4.0 are edge devices. As an evolution of Programmable Logic Controllers (PLC), they are frequently used for retrofitting, e.g., equipping legacy manufacturing machines with recent communication protocols. Moreover, modern edge devices combine hard real-time functions connected to the manufacturing machines with soft/non real-time IT capabilities. Some recent edge devices even ship with modern hardware accelerators, such as Graphic Processing Units (GPUs) or Tensor Processing Units (TPUs), which are often beneficial for AI calculations. While edge devices allow for offloading IT functionality close to production machines, e.g., to operate AI at low latency, they also significantly increase the management and deployment complexity in Industry 4.0 setups by emphasizing distributed on-premise computing.

To support companies in managing this complexity, several software platforms for Cyber-Physical Production Systems (CPPS) or Industrial Internet of Things (IIoT/IoT) applications are available, e.g., Siemens MindSphere or PTC ThingWorx. As we discussed in [3], these platforms significantly differ in their capabilities, in particular with respect to AI, edge offloading or cloud usage. Moreover, they often fall short in providing capabilities for consistent (distributed) system customization, one key capability to cope with the complexity, but also in data protection or data/service sharing for ecosystem setups.

In the funded project IIP-Ecosphere, we are researching concepts for easy-to-use AI in the manufacturing domain. The overall mission of IIP-Ecosphere is to create an ecosystem of involved stakeholders for the mutual transfer of experience and knowledge. For demonstrating the approaches, the partners develop a prototypical IIoT platform. On the one side, the requirements for such a platform must reflect the scientific goals and pave the way for experimenting with and demonstrating of novel approaches. On the other side, such a platform must also be interesting for industrial stakeholders and support production requirements. Thus, an elicitation of platform requirements needs to be carefully balanced.

Our main questions are 1) how to collect and combine scientific and industrial requirements in an Industry 4.0 context and 2) can different views on the requirements be used to improve their mutual completeness? As contributions we present a pragmatic combination of scientific methods, e.g., surveys, with requirements elicitation techniques in the context of an industrial reference process for systems design. This involves the creation two complementing views, a usage, as well as a functional/quality view on the requirements and allows for a more encompassing requirements collection, but also a discussion of mutual influences. We provide insights into elicited requirements and experiences that we made.

In summary, we collected 67 usage view activities/scenarios and 141 top-level requirements further detailed by 179 sub-requirements. These requirements characterize the (prioritized) desires for an AI-enabled Industry 4.0 platform. Intentionally,

 ISBN: 978-1-61208-946-1

we were open to requirements that will probably not be realized during the lifetime of IIP-Ecosphere in order to provide inspiration for future works. At the point of writing, about 35% of the requirements have been realized in a prototypical open source platform and several of the identified concepts are being taken up by a standardization initiative for edge devices in Industry 4.0. Moreover, some industrial IIP-Ecosphere partners adopted our integrated requirements approach to improve their internal software development processes.

This paper is structured as follows: In Section II, we provide a brief overview of the IIP-Ecosphere project. In Section III, we introduce our approach for requirements collection and discuss results from that approach in the following sections, i.e., on a detailed platform survey in Section IV and for the requirements collection with two views in Section V. In Section VI, we discuss related work and in Section VII we conclude this paper and outline future work.

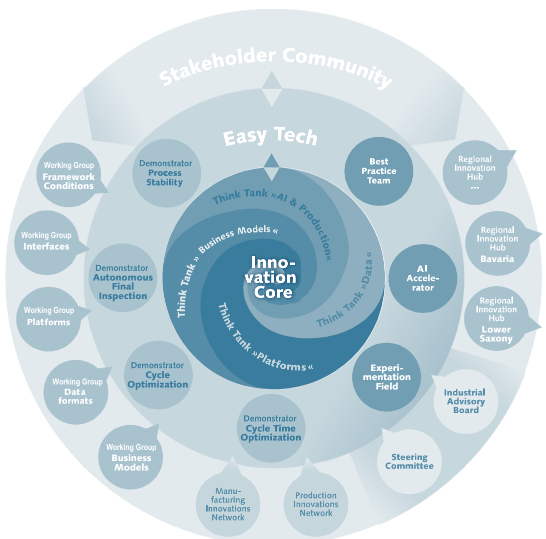

Figure 1. IIP-Ecosphere project structure.

II. IIP-Ecosphere Project Overview

Our work takes place in the context of the IIP-Ecosphere project, which is funded by the German ministry for Economics and Energy in its AI innovation competition program. IIP-Ecosphere aims at achieving an innovative leap in the field of industrial production exploiting networked, intelligent, autonomous system capabilities to increase productivity, flexibility, robustness and efficiency of Industry 4.0. The goal is to build a novel ecosystem of humans (through companies and organizations), software, machines and products with a specific focus on mutual experience and knowledge transfer.

To achieve this, the activities in IIP-Ecosphere are structured in three layers, as illustrated in Figure 1. The *Innovation Core* is at the heart of the ecosystem and is constituted by four so called think tanks performing research on core topics, such as platforms, AI, business models and data. The *Easy Tech* layer aims at demonstrating the research results and transferring them into industrial practice, in particular through the *AI Accelerator*, which works, e.g., on a public catalog of AI solutions and on generalized, (re-)usable AI services for manufacturing. Finally, the *Stakeholder Community* conducts activities for external parties, e.g., workshops on the core topics or linking of linking start-ups, SMEs, large companies and multiplicators with the project (*Regional Innovation Hubs*). This paper is based on joint activities of the platform think tank, the AI accelerator and the demonstrators. After the end of the project's lifetime, the created community/ecosystem shall continue the project activities on its own.

One core activity in IIP-Ecosphere is the realization of a virtual platform that connects existing devices and factory installations in a vendor-independent manner. A *virtual platform* [4] takes up functionality and services of existing, already installed protocols and platforms, integrates them and offers additional services on top of these. In IIP-Ecosphere, we aim at enabling intelligent manufacturing applications based on an open set of re-usable AI and platform services. These services shall be flexibly distributed to available resources, such as edge devices, on-premise servers or clouds. The service distribution shall be determined by the platform before starting an application, but also during run-time, i.e., in a self-adaptive manner. As requested by the funding scheme, IIP-Ecosphere strives for concepts and methods to achieve/increase vendor-neutrality, interoperability and flexible uptake of Industry 4.0 related standards, e.g., Message Queuing Telemetry Transport (MQTT) [5], or Open Platform Communications Unified Architecture (OPC UA) [6].

III. Approach

The realization of such a platform is not only a technical endeavor that commands the application of software engineering methods, such as requirements engineering or architectural design. It forms a data-driven system and, thus, faces challenges that are, e.g., discussed in [7]. Particular challenges are highly interdisciplinary teamwork (production, AI, data science, software engineering, economics) including researchers and practitioners, but also volatile and unclear requirements due to explorative AI and data science processes.

As stated above, we head for a research-integrated requirements collection, which is based on relevant standards/approaches for Industry 4.0 and IIoT. For system development, the German Standardization Roadmap Industry 4.0 [8] advocates the Industrial Internet Reference Architecture (IIRA) [9], in particular the so-called 'Industrial Internet Viewpoints'. Figure 2 a) illustrates these interlinked viewpoints, consisting of a *business view* (roles attributed with business interests), an *usage view* (a use case collection for all involved roles and system entities), a *functional view* (domain decomposition of system functions) and an *implementation view* (detailed architecture). This approach is also used in relevant inputs for our work, particularly in an international effort to standardize edge computing in manufacturing [10]. However, like several other works [11]–[13], the IIRA approach focuses on the technical side, neglecting research demands.

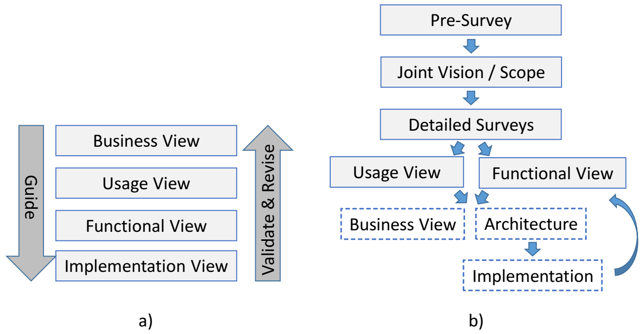

Figure 2. Steps towards requirements: a) IIRA [9] b) our approach.

For our requirements collection, we adopted the IIRA usage and functional views in Figure 2 b) as follows:

- Start with an open-minded **pre-survey**: We conducted surveys on research literature for IIoT platforms and on economically predominant IIoT platforms. As result, we identified (research-)gaps in dynamic and adaptive deployment, semantic data integration, security, and consistent customization/configurability (in the sense of variability modeling in software product lines [14]).
- Create a **joint vision**: Based on the pre-surveys, we identified further (research-)relevant topics and integrated them into a joint vision. One topic is to explore the upcoming Asset Administration Shell (AAS) [15] standard, which aims at interoperable modeling of Industry 4.0 ”assets”, i.e., products, machines or digital twins, similar to the “Smart Manufacturing Profile” concept in the US. From a software perspective, AAS can be viewed as (distributed) functional interfaces allowing for transparent remote access [16]. One aim is to identify benefits and limitations of AAS, e.g., platform interfaces can efficiently be realized by AAS. Further platform challenges target transparent mechanisms for data privacy, secure data sharing (along the lines of the International Data Spaces Association (IDSA) [17]) or the optimization of code deployment to computational resources.
- Stabilize the vision by **detailed surveys**, i.e., assure the gaps and identify supporting arguments for the vision through focused surveys. In Section IV, we will report on a survey of IIoT platforms, while an accompanying systematic literature review is out of scope here.
- Create a **usage** and a **functional** view: Using the vision as scope, elicit the requirements in terms of the two views so that they can complement each other. In our case, two teams created the views and performed a comparison of the results to assess and improve the comprehensiveness of the requirements collection. We will report our results for both views, the identified similarities and gaps, as well as our experiences in Section V.

Our results act as input for further works, e.g., the IIP-Ecosphere think tank “Business Models” uses our usage view to derive a business view for the platform and ultimately for the ecosystem. Further, the technical partners design an architecture and create an implementation of the platform based on the collected research-integrated requirements. In turn, this will act as a basis for the think thanks and demonstrators to demonstrate their results in an integrating environment.

It is important to emphasize that the collected requirements are so far based on the input of the IIP-Ecosphere partners. Ongoing work with the stakeholder community may lead to additional input and a refinement of the existing views. This input may be taken up in an iterative manner or induce requirements that document future work for the community.

IV. Survey of current IIoT Platforms

To stabilize the joint vision, we performed a survey of current IIoT platforms [3]. We selected 21 platforms, among them 9 platforms due to a competitive stakeholder analysis (including AWS IoT, PTC ThingWorx, SAP Leonardo or Siemens MindSphere), as well as 12 further platforms of particular interest to the project (such as Adamos, Deviceinsight Centersight, or Software AG Cumolocity). Based on a pre-survey and the joint vision, we defined 16 analysis dimensions including (re-usable) AI, Edge/IoT/cloud capabilities, uptake of standards, security, data privacy, ecosystem building and systematic configurability. We systematically analyzed the platforms along these dimensions based on vendor material and web pages in the period from June to August 2020.

- Although stated as relevant to almost all platforms, only 77% detail their **AI capabilities**. 48% enable customizations of the AI capabilities, while only 14% support user-defined orchestration or third party AI functionality.
- 95% of the platforms offer some form of **cloud** integration, which is frequently used to argue the scalability of the platform. Although a (mandatory) cloud integration is sometimes perceived by customers as an adoption barrier [18], only 19% offer an optional cloud integration, and only 24% support an **on-premise** installation.
- 85% support **edge** devices, but the functionality is rather diverse, ranging from data storage (67%) to customer-specific deployments (29%). 33% support AI on edge devices, however, this is currently often limited to functionality shipped with the platform. 38% of the platforms rely on container technology (usually Docker [19]) and 4 platforms (19%) utilize containers for edge deployment.
- 57% are characterized as (soft-)**real-time** capable. This roughly correlates with the edge findings. 76% employ some form of data stream or complex event processing, partially offering query languages, "low code" or "no code" environments to customize the data processing.
- Usually, the platforms offer extensive support for modern and legacy **protocols**, as well as (secure) device management. More recent approaches like OPC UA are used rarely. Most of the platforms offer some vendor-specific (REST) **interfaces**, while none of the platforms seems to uptake recent Industry 4.0 interfacing works like AAS [15] or OPC UA companion specs.

 ISBN: 978-1-61208-946-1

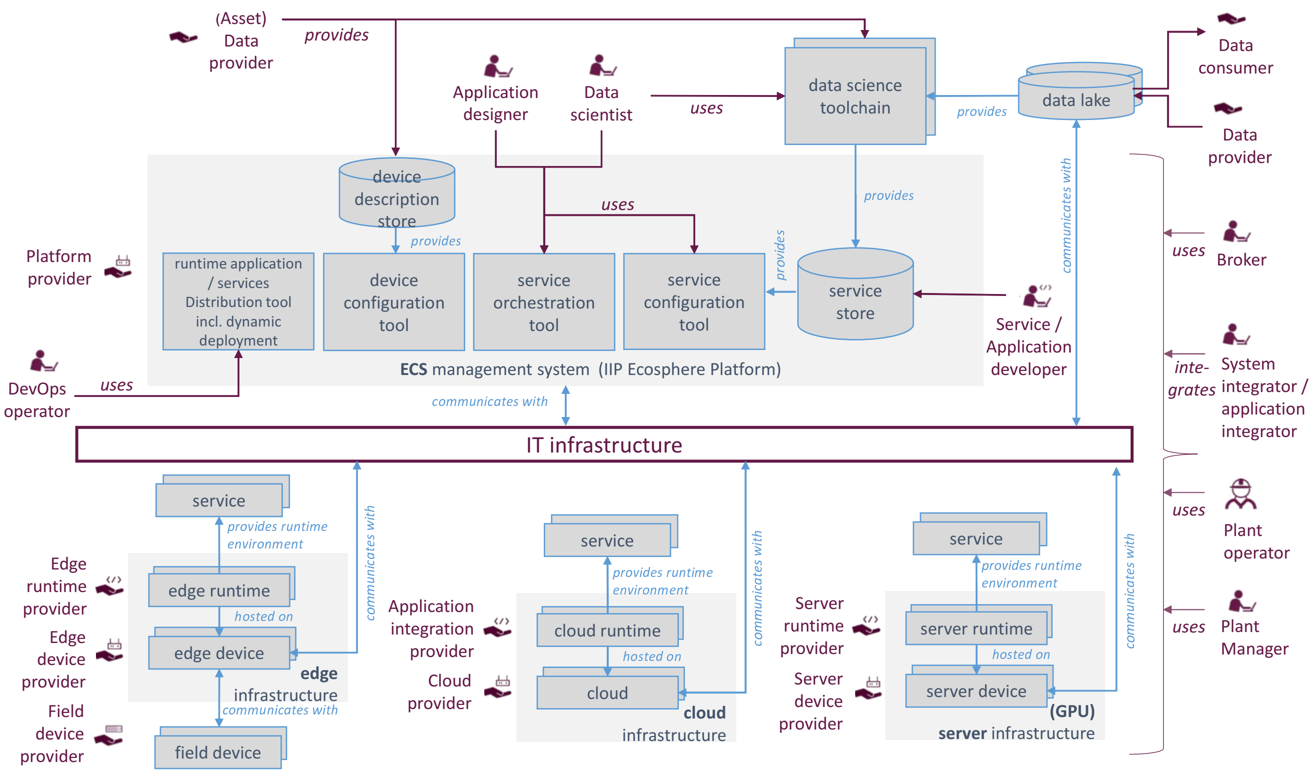

Figure 3. IIP-Ecosphere System under Consideration (SuC).

- **Security** and **data protection** seem to be essential for all surveyed platforms, in particular for cloud integration. 86% describe employed authorization measures, 71% allow limiting the data retention, but only 48% implement mechanisms to control processing of personal data.
- 81% of the platforms appear to be **customizable**, e.g, 62% of the surveyed platforms allow for external (AI) components. However, the utilized mechanisms, e.g., to ensure a consistent platform configuration, remain unclear despite the fact that customization approaches for the manufacturing/CPPS domain do exist, e.g., [20].
- Openness and customization often correlate with platform **ecosystems** [4]. Usually, the platforms focus on developer and community support, while only some platforms build up an ecosystem around their own platform (19%).

With the advent of AI, the demand for real-time processing and flexible deployments of (customer-defined) AI methods will become more prevalent. This coincides with demands for flexible offloading including edge devices for latency reduction and cloud capabilities. However, issues in standardization, openness, interoperability among platforms, security and data protection/privacy impact this trend, as well as the user's freedom of choice. We used these results to confirm the gaps/topics identified for the joint platform vision (cf. Section III) and as a scope for the subsequent requirements collection.

V. Requirements Collection

We now detail the requirements collection for the IIP-Ecosphere platform, the results and experiences that we made.

The requirements collection was conducted by two distinct teams. The input was mostly elicited through document analysis (relevant papers, standards and documents also attesting prior work, as well as the IIP-Ecosphere grant agreement as described in [21]) and interactive workshops, with stakeholders from research, industry and multipliers with backgrounds in AI (research, application), industrial production, factory construction, device supply, software engineering in individual cases also with experience in requirements collection. The workshops introduced the vision/scope based on the findings from our review of selected IIoT platforms [3], gave an explanation of the respective approach to requirements collection, and typically led to many interactions and lively discussions. The results of these interactions were scripted, summarized in a document and reviewed by the participants.

A. *Usage view*

An IIRA usage view consists of an initial architecture, the "System under Consideration" (SuC), a definition of the used *entities*, the interacting *roles*, as well as *activities* on when/how the *roles* interact with the *entities*. Activities can be specified in terms of a template with a sequence enumerating the interactions (similar to use cases). The edge configuration usage view [10] provided a good basis for our work, but it does

 ISBN: 978-1-61208-946-1

not cover all relevant topics for IIP-Ecosphere, particularly deployment to on-premise/cloud resources or AI activities. For these topics we organized focus workshops, where the participants discussed existing/new roles, entities and activities. Finally, we integrated the collected information into the usage view in [22] with the following results:

- A significantly **extended SuC**, as illustrated in Figure 3. The entities are colored in gray, the roles in purple, and the interactions between the SuC and roles are drawn as purple arrows. Areas indicating systems of the IIP-Ecosphere platform, such as the ECS (Edge, Cloud and Server) management system or the ECS devices/infrastructures, are shaded in light gray, covering all entities that belong to the indicated systems. Data flows between entities are drawn as light blue arrows. The underlying IT Infrastructure of the SuC, connecting the ECS devices/infrastructures with the IIP-Ecosphere platform is depicted as a white box with a purple outline, as it is neither entity nor role in the SuC.
- As in [10], **field devices**, such as sensors, actuators or (parts of) manufacturing machines are only connected to edge devices. In contrast to [10] where entire applications are deployed to edge devices, the IIP-Ecosphere platform shall focus on **applications** that are composed of orchestrated services that can be distributed across ECS devices.
- Applications and services are specified in a **configuration model** (not shown in the SuC) that allows for creating the runtimes of applications and services for the deployment into containers and for determining optimizations or incompatibilities in the orchestration of services. Means for validating the consistency of the model shall be provided and integrated with the user interactions. Services can be added to/updated in a service store as needed.
- Each ECS device runs a **device abstraction** (ECS runtime) being responsible for executing the services/containers and for reporting their runtime measurements. The platform decides about the deployment, dynamically composes service containers for the target resource, and adapts the distribution runtime. For developing applications and services, the stakeholders indicated that **pre-deployment testing or simulation** of new or updated services is highly desirable.
- All resources and services provide a **self-description** in terms of an AAS [15] information model and communicate only via **Industry 4.0 protocols** to foster interoperability, but also to explore limitations.
- The elicited **AI and data science activities** form an orthogonal space. The stakeholders contributed activities for data exploration, AI model design/testing and the integration of external data science toolchains. In the context of developing new AI services, the stakeholders also expressed the need for the ability to use **pre-deployment testing** of new and updated AI services/models, which are key elements in the continuous development and operation of machine learning (MLOPs, [23]) of applications and services within the IIP-Ecosphere platform.
- AI methods typically operate on models that may incur **data protection, IPR or further legal issues**. Some issues may be addressed by limiting deployment targets, e.g., through the exclusion of certain cloud spaces. Issues, such as data protection could be addressed by modifying the data close to the source, e.g., through anonymization or pseudonymization. However, the impact of such modifications on AI and further data processing is currently unclear and shall be researched using the platform.

In summary, as illustrated in Figure 4, the usage view for the IIP-Ecosphere platform consists of 18 entities, 19 roles, as well as 43 deployment and 24 AI activities (as opposed to [10] with 5 entities, 7 roles and 27 activities). Although the joint vision focused the discussions to a certain system scope and one might expect that this also limited the contributions of the participants, several results were creative and surprising to us. We will detail some examples below.

For example, the data scientists argued about alternatives on how to integrate a data science toolchain into the platform. An initial idea on a tight integration was rejected in favor of a loose integration, i.e., the toolchain shall remain flexible so that a data scientist may use his/her favorite tools while the platform supports the process in terms of provisioned resources, access to experimental and life data, as well as available (AI) services. This insight let to 10 activities specifically on 'activities for model training and evaluation', which cover many of the 9 stages of the machine learning workflow in [24], some are exemplified in Figure 4. Furthermore, the aspect of continuous delivery of AI services following MLOps concepts was emphasized by the AI experts. This led to activities like 'continuous application of AI models on new data' or 're-calibration of AI model parameters' shown in Figure 4. Moreover, the stakeholders requested capabilities to measure the accuracy of productive or simulated AI models to observe the quality of predictions and to early on detect model-drift, e.g., the loss of accuracy due to slow changes in the application environment. This induced 5 activities on 'analysis and prediction of performance and accuracy'.

During the workshops, the industry experts expressed the necessity to provide simulation- and testing-capabilities to allow for simulation-driven development of applications and services. A key approach that was identified here is the development, simulation and monitoring of applications and services but also of ECS devices, based on digital twins. Similar to MLOps, these activities target DevOps [25] capabilities, for example, allowing for pre-deployment testing of any application, service or ECS device. We represented this desire in terms of 8 'activities for (distributed) applications'.

As described, the presented results focus on service deployment and AI activities. Initially, we planned to explore also further topics, such as data sharing or data privacy. However, we also experienced that interactively creating a usage view is a significant effort. Thus, in particular for a research-integrating usage view, it is important to focus on the most important topics first. It is noteworthy that, as outlined above, a

- **Entities** (18)
- **Roles** (19)
  - *Edge device provider*
  - *Data scientist*
- **Activities for ECS management** (7)
  - Adding Entities (4)
  - Removing Entities (4)
  - Provision of Entities (8)
  - Provision of service and application template (4)
  - Service configuration and orchestration (6)
  - Setting up operational configurations (9)
  - Activities for (distributed) applications (8)
    - *Simulating the integration*
    - *Simulating the deployment*
    - *Visualizing the results*
- **Activities for AI services and processes** (5)
  - Activities for data exploration (5)
  - Activities for model training and evaluation (10)
    - *Training of AI models*
    - *Provision of intermediary model results*
    - *Continuous application of a model on new data*
    - *Re-calibration of model parameters*
  - Use of AI applications/services (2)
  - Analysis/Prediction of performance and accuracy (5)
    - *Provision of metrics for an application/service*
    - *Analysis of metadata to detect deviation/model drift*
  - Using AI services/applications manually, offline (2)

Figure 4. Usage view overview with sections, (number of) contained roles / activities, and example roles / activities (in italics).

significant number of capabilities needed for the platform and subsequent activities that enable theses, were identified by the close communication and interaction between the AI experts, who were focused on research aspects of the IIP-Ecosphere platform and the Industry experts with their focus lying on the technical aspects of the platform. Hence, the integration of the scientific and the industrial view in these two groups yielded deep and very valuable results for the elicitation and formulation of the IIP-Ecosphere platform requirements. It is also worth to mention that during the cooperation of both expert groups a significant amount of "mutual understanding" was established, clarifying for both groups of experts specific vocabulary, views and motivations present in the two groups and thus enabling a productive discourse and collaboration.

B. Functional View

A second team collected required platform functions, as well as quality requirements, initially independent of the usage view activities. We performed a requirements collection combining the recording of ideas and desires mentioned in specific discussions with structured approaches, such as interviews or questionnaires. In particular, within the consortium we conducted a requirements questionnaire with 8 questions driven by the joint vision, ranging from a summary of the planned applications over envisioned AI methods, relevant data protection measures up to imaginable run-time changes for self-adaptation.

We documented the requirements in terms of phrase templates [26], i.e., based on a *template sentence* indicating the acting role, the required functionality and the prioritization of the requirement. All requirements were reviewed by the stakeholders, categorized, prioritized (must/should/can) - with more emphasis on scientific goals, required basis functionality and the grant agreement of IIP-Ecosphere - annotated with their source, and, if needed, detailed by an explaining text.

Ultimately, we compared the usage view with the functional view. While more than two third of the topics do occur in both views, we also identified differences. We found entire topics in the usage view that the stakeholders did not touch at all in the requirements discussions, e.g., the pre-deployment simulation. Moreover, the interaction steps in the usage view activities pointed us to details that were not covered by the requirements, e.g., how IIoT applications shall be managed. For the opposite direction, we found, e.g., that run-time adaptation was treated in the questionnaire as an interesting feature, which was also viewed with caution, i.e., some stakeholders requested explicit human approvals rather than autonomously changing a deployment (or, similarly, a re-trained AI model).

In summary, we elicited 141 top-level and 179 refining sub-requirements as documented in [21]. 17% of all requirements were added due to the comparison of the two requirements views. Figure 5 illustrates the requirements categories that we identified along with the number of contained top-level and sub-requirements, as well as selected example requirements (without explaining text). About 16% of all requirements target quality, among them 7 on data processing, e.g., on the expected data frequency and volume. The largest group of quality requirements focuses on security and data protection. The main sources are the think tanks (41%), the IIP-Ecosphere demonstrators (20%) and the grant agreement (12%). Further sources are standards, the platform survey from Section IV and the comparison with the usage view.

Although our set of functional requirements is rather detailed, we are aware that it is potentially incomplete. On the one side, the IIP-Ecosphere platform forms a data-driven system and, as mentioned above, requirements in such systems are known to be volatile, unclear or incomprehensive due to the explorative nature of data science and AI processes [7]. On the other side, resource limitations in this research project prevented us from conducting further/deeper usage view and

- **General Requirements** (12, 15)
- **Connectors and Connections** (10, 20)
- **Heterogeneous, dynamic Deployment** (15, 24)
  - *R24. Resource properties/functions must be described as AAS.*
  - *R26. Platform must support on-premise deployment.*
- **Security** (7, 6)
- **Data Protection** (24, 8)
- **Central Storage Services** (10, 21)
- **Data Sharing** (4, 8)
- **Data Integration** (10, 0)
- **Configurability** (9, 9)
  - *R94. Platform must support automatic configuration validation.*
  - *R96. Configuration must include optional/alternative components/services.*
- **Optimized / Adaptive Deployment** (8, 9)
- **AI (Service) Toolkit** (10, 27)
  - *R110. The AI toolkit must define interfaces for AI components in industrial production.*
  - *R111. The AI toolkit must be extensible.*
- **Adaptive Service Selection** (7, 6)
- **Virtualization** (4, 0)
- **Application Support** (11, 26)

Figure 5. Requirements overview with categorization topics, number of top-/sub-requiements and selected example requirements (in italics).

requirements workshops, e.g., on data sharing. Thus, in both views, particular topics/requirements could still be missing. However, the described results allow for discussing the effects of an interactive multi-view requirements collection. Moreover, the collected requirements are sufficient to incrementally realize the IIP-Ecosphere platform and, if needed, to elicit missing details during platform realization in agile manner.

C. Experiences

The two teams recorded the experience that they made during the elicitation that we briefly summarize in this section.

Both views are based on templates (activities and phrases) that seemed to give a certain form of guidance to the participants. In comparison, the usage view workshops seemed to have allowed for more creativity, i.e., the participants tended to talk more freely about desirable system interactions or known limitations. The workshops also allowed for more interactions, such as agreements among the participants as stated in Section V-A. However, this impression may be biased by the first workshop, where the participants were asked to name any missing topic. Some of the topics were taken up by later workshops where then mostly 'experts' participated, which allowed also more quiet persons to participate more actively.

The discussions on the functional view were focused on capabilities for developing applications on top of the platform and, thus, more technical. Here the stakeholders did not interact so lifely, which might be one reason why we missed application-related topics. Moreover, we noticed that different persons participated in the usage view and in the requirements workshops. In the latter, the participants seemed to have more technical background, probably as the workshop invitation asked for contributions to functional requirements for the platform. Yet, the functional view also revealed interesting aspects, e.g., as mentioned above, (different levels) of adaptation approvals or the need for explainable adaptation decision making. This may be biased by our questionnaire, where we explicitly asked for these topics and the participants could overthink their answers or discuss them with their team.

We also experienced that research-integrated requirements do not come for free. Questions like "Why do we need this?" or "Isn't that too risky?" for certain research topics, e.g., for self-adaptive capabilities, arise and must even be defended against more practical/industrial requirements.

VI. Related Work

We now review briefly work related to our core topics, i.e., surveys and requirements collections for IoT platforms.

Various comparisons and *surveys of IoT platforms* are published. As stated by Mijuskovic at al. [27], this is often done for a specific set of criteria lacking a sound comparison framework. Moreover, comparisons are typically based on a selection of platforms as the market is rather dynamic and encompasses hundreds of platform vendors [28], i.e., typical numbers of platforms are 11 in [29], 13 in [30], 20 in [31], 24 in [28] or 26 in [32]. Often, such surveys are based on vendor material, while in [28] interviews with vendor representatives were used. Regarding comparison criteria, the topics are frequently device management [28] [30] [31], fog/edge/cloud deployment [29] [30] [32], connectivity/protocols [28] [29] [30] [31] [32], security [28] [29] [31], data management [28] [29] [30] [32], data analytics [28] [31] [32], visualizations/UI [28] [31] [32], application development [28] [29] [30] [32], system/service management [30] [32], or licensing/payment [29] [30]. In contrast, in our survey we also analyzed AI capabilities, edge usage, ecosystem building, data protection and consistent configurability and used that survey as a basis for our research-integrated requirements approach.

There is also a body of work on requirements management for IoT/CPPS platforms or ecosystems, e.g., [11]–[13]. However, we do not aim at proposing a completely new requirements approach rather than performing a *requirements collection* for IoT/CPPS systems while balancing scientific and practical interests. Many technical publications motivate their work with a focused set of requirements, while overview work with collections of platform requirements is less common. Among those, we identified the following topics for functional requirements: device/resource/distribution management [33] with heterogeneous deployment [27], communication/networking [27] [34], data (base) management [27] [33] [34], data processing [27] [33], data analytics including AI [27], monitoring [27] service management [34], security/privacy [27] [33] [35] [36], or visualization [27] [33]. Moreover, we found non-functional topics, such as scalability [33] [35], performance [27] [33] [34], standardization/interoperability [33]–[35], development support [27] [33] and even self-adaptation [33]. In contrast to our work, the cited publications typically focus on a single (functional/quality requirements) view, i.e., do neither take the scientific site nor interactions of multiple views into account.

VII. Conclusion

IIoT, CPPS and Industry 4.0 platforms form the software foundation of complex manufacturing systems. The introduction of Artificial Intelligence into such systems will enable new opportunities, but further increases the complexity and causes challenges for all involved disciplines. Eliciting requirements for future platforms is not trivial, in particular if scientific and industrial interests must be balanced and integrated.

In this paper, we reported on a pragmatic approach to perform a requirements collection of a platform that shall demonstrate research approaches in an upcoming Industry 4.0 ecosystem. Driven by pre-studies, we used a joint vision as scope for the further steps, a surveying phase and a requirements collection phase. For the surveys, we reported on an overview of 21 recent Industry 4.0 platforms that helped us to identify gaps and to stabilize the vision. The vision then guided an intensive requirements collection for two perspectives, a usage view and a functional/quality view, which, in summary, integrates research and industrial interests. We discussed our experiences with such a requirements elicitation, in particular that different views can successfully complement each other.

The joint work, in particular on documenting the results, helped the involved partners to clarify and synchronize their view on the system to be built, e.g., the terminology or the needed components. Based on these experiences, first companies in IIP-Ecosphere started applying such a requirements elicitation approach as part of their own activities. Moreover, concepts and ideas on service-based Industry 4.0 platforms as outlined in the usage view [22] were fed back to the originating Labs Network Industrie 4.0 (LNI 4.0) organization and at the time of writing are being integrated into a revised version of [10].

Current and future work is on developing the IIP-Ecosphere platform based on both requirements documents, including incremental architecture design or integration of research and industrial approaches. At the time of writing, about a third of all platform requirements have been realized and validated. We also plan for evaluations of the platform approaches in terms of industrial use cases.

ACKNOWLEDGMENT

The IIP-Ecosphere project (https://www.iip-ecosphere.de/) is partially supported by the German Ministry of Economics and Energy (BMWi) under the grant numbers 01MK20006C and 01MK20006D. The authors would like to thank the partners who participated in and contributed to our requirements workshops.

REFERENCES

[1] J. Delsing, "Local Cloud Internet of Things Automation," *IEEE Industrial Electronics Magazine*, pp. 8–21, 2017.
[2] R. Wöstmann *et al.*, "Conception of a Reference Architecture for Machine Learning in the Process Industry," in *Intl. Conference on Big Data*, 2020, pp. 1726–1735.
[3] C. Sauer, H. Eichelberger, A. S. Ahmadian, A. Dewes, and J. Jürjens, "Current Industry 4.0 Platforms - An Overview," retrieved 03-2022. [Online]. Available: https://www.iip-ecosphere.eu/wp-content/uploads/2021/02/IIP-2020_001-en.pdf
[4] K. Schmid, H. Eichelberger, and C. Kröher, "Domain-Oriented Customization of Service Platforms: Combining Product Line Engineering and Service-Oriented Computing," *Journal of Universal Computing*, vol. 19, no. 2, pp. 233–253, 2013.
[5] "Website of the Organization for the Advancement of Structured Information Standards on MQTT," retrieved 03-2022. [Online]. Available: https://mqtt.org/
[6] "Website of the OPC UA foundation," retrieved 03-2022. [Online]. Available: https://opcfoundation.org/
[7] O. Hummel, H. Eichelberger, A. Giloj, D. Werle, and K. Schmid, "A Collection of Software Engineering Challenges for Big Data System Development," in *Proceedings Euromicro Conference on Software Engineering and Advanced Applications (SEAA'18)*, 2018, pp. 362–369.
[8] H. Bedenbender *et al.*, "Industrie 4.0 Plug-and-Produce for Adaptable Factories: Example Use Case Definition, Models, and Implementation," 2017. [Online]. Available: https://www.zvei.org/fileadmin/user_upload/Presse_und_Medien/Publikationen/2017/Juni/Industrie_4.0_Plug_and_produce/Industrie-4.0-_Plug-and-Produce-zvei.pdf
[9] Industrial Internet Consortium, "The Industrial Internet Reference Architecture," retrieved 03-2022. [Online]. Available: https://www.iiconsortium.org/pdf/IIRA-v1.9.pdf
[10] A. Graf, U. Löwen, M. Rentschler, B. Vojanec, and T. Roehrl, "LNI 4.0 Testbed Edge Configuration - Usage View," retrieved 03-2022. [Online]. Available: https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/LNI4.0-Testbed-Edge-Configuration_UsageViewEN.html
[11] P. Loucopoulos, E. Kavakli, and N. Chechina, "Requirements Engineering for Cyber Physical Production Systems," in *Advanced Information Systems Engineering (CAiSE'19)*, 2019, pp. 276–291.
[12] S. Kaleem *et al.*, "A Review on Requirements Engineering for Internet of Things (IoT) Applications," in *HCT Information Technology Trends (ITT'19)*, 2019, pp. 269–275.
[13] D. Silva, T. G. Gonçalves, and A. R. C. da Rocha, "A Requirements Engineering Process for IoT Systems," in *Brazilian Symposium on Software Quality (SBQS'19)*, 2019, pp. 204–209.
[14] F. van der Linden, K. Schmid, and E. Rommes, *Software Product Lines in Action*. Springer, 2007.
[15] S. Bader *et al.*, "Details of the Asset Administration Shell," retrieved 03-2022. [Online]. Available: https://www.plattform-i40.de/PI40/Redaktion/DE/Downloads/Publikation/Details_of_the_Asset_Administration_Shell_Part1_V3.html
[16] M. G. Casado and H. Eichelberger, "Industry 4.0 Resource Monitoring - Experiences with Micrometer and Asset Administration Shells," in *Symposium on Software Performance (SSP '21)*, 2021.
[17] "Website of the International Data Space Association," retrieved 03-2022. [Online]. Available: https://internationaldataspaces.org/idsa-industrial-community/
[18] C. Esposito, A. Castiglione, B. Martini, and K. R. Choo, "Cloud manufacturing: Security, privacy, and forensic concerns," *IEEE Cloud Comput.*, vol. 3, no. 4, pp. 16–22, 2016.
[19] "Website of Docker Inc." retrieved 03-2022. [Online]. Available: https://www.docker.com/
[20] K. Meixner, R. Rabiser, and S. Biffl, "Towards Modeling Variability of Products, Processes and Resources in Cyber-Physical Production Systems Engineering," in *Intl. Systems and Software Product Line Conference (SPLC'19) - Vol. B*, 2019, pp. 49–56.
[21] H. Eichelberger *et al.*, "IIP-Ecosphere Platform Requirements (Functional and Quality View)," retrieved 03-2022. [Online]. Available: https://www.iip-ecosphere.eu/wp-content/uploads/2021/03/IIP-2021_002-en.pdf
[22] C. Sauer, H. Stichweh, and H. Eichelberger, "IIP-Ecosphere Platform Requirements (Usage View)," retrieved 03-2022. [Online]. Available: https://www.iip-ecosphere.eu/ergebnisse/
[23] S. Mäkinen, H. Skogström, E. Laaksonen, and T. Mikkonen, "Who needs mlops: What data scientists seek to accomplish and how can mlops help?" *arXiv preprint arXiv:2103.08942*, 2021.
[24] S. Amershi *et al.*, "Software engineering for machine learning: A case study," in *International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 2019, pp. 291–300.
[25] R. Jabbari, N. bin Ali, K. Petersen, and B. Tanveer, "What is DevOps? A systematic mapping study on definitions and practices," in *Scientific Workshop Proceedings of XP2016*, 2016, pp. 1–11.
[26] K. Pohl and C. Rupp, *Requirements Engineering Fundamentals*. Rocky Nook, 2011.
[27] A. Mijuskovic, I. Ullah, R. Bemthuis, N. Meratnia, and P. Havinga, "Comparing Apples and Oranges in IoT context: A deep dive into methods for comparing IoT platforms," *IEEE Internet of Things Journal*, pp. 1797–1816, 2020.
[28] T. Krause *et al.*, "IT-Plattformen für das Internet der Dinge (IoT)," 2017. [Online]. Available: http://publica.fraunhofer.de/documents/N-470532.html
[29] P. Agarwal and M. Alam, "Investigating IoT Middleware Platforms for Smart Application Development," in *Smart Cities - Opportunities and Challenges*, ser. Lect. Notes Civ. Eng. Springer, 2020, vol. 58.
[30] M. Asemani, F. Abdollahei, and F. Jabbari, "Understanding IoT Platforms: Towards a comprehensive definition and main characteristic description," in *Intl. Conference on Web Research (ICWR'19)*, 2019, pp. 172–177.
[31] H. Hejazi, H. Rajab, T. Cinkler, and L. Lengyel, "Survey of platforms for massive IoT," in *Future IoT Technologies*, 2018, pp. 1–8.
[32] P. P. Ray, "A survey of IoT cloud platforms," *Future Computing and Informatics Journal*, vol. 1, no. 1, pp. 35 – 46, 2016.
[33] P. Agarwal and M. Alam, "Open Service Platforms for IoT," in *Internet of Things (IoT): Concepts and Applications*, 2020.
[34] E. Patti and A. Acquaviva, "IoT platform for Smart Cities: requirements and implementation case studies," in *Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI'16)*, 2016.
[35] P. Bhuyan and A. Ray, "IoT Service Platform," in *Interoperability in IoT for Smart Systems*. CRC Press, 2021, pp. 91–97.
[36] S. Oh and Y. Kim, "Security Requirements Analysis for the IoT," in *International Conference on Platform Technology and Service (PlatCon'17)*, 2017, pp. 1–6.