

Guidelines for Artifacts to Support Industry-Relevant Research on Self-Adaptation

Danny Weyns^{1,2}, Ilias Gerostathopoulos³, Barbora Buhnova⁴, Nicolás Cardozo⁵, Emilia Cioroai⁶, Ivana Dusparic⁷, Lars Grunske⁸, Pooyan Jamshidi⁹, Christine Julien¹⁰, Judith Michael¹¹, Gabriel Moreno¹², Shiva Nejati¹³, Patrizio Pelliccione¹⁴, Federico Quin¹, Genaina Rodrigues¹⁵, Bradley Schmerl¹², Marco Vieira¹⁶, Thomas Vogel^{8,17}, Rebekka Wohlrab¹²

¹Katholieke Universiteit Leuven; ²Linnaeus University; ³Vrije Universiteit Amsterdam; ⁴Masaryk University; ⁵Universidad de los Andes; ⁶Fraunhofer IESE; ⁷Trinity College Dublin; ⁸Humboldt-Universität zu Berlin; ⁹University of South Carolina; ¹⁰University of Texas Austin; ¹¹RWTH Aachen University; ¹²Carnegie Mellon University; ¹³University of Ottawa; ¹⁴Gran Sasso Science Institute; ¹⁵Universidade de Brasília; ¹⁶University of Coimbra; ¹⁷Paderborn University
danny.weyns@kuleuven.be, ilias.gerostathopoulos@vu.nl

ABSTRACT

Artifacts support evaluating new research results and help comparing them with the state of the art in a field of interest. Over the past years, several artifacts have been introduced to support research in the field of self-adaptive systems. While these artifacts have shown their value, it is not clear to what extent these artifacts support research on problems in self-adaptation that are relevant to industry. This paper provides a set of guidelines for artifacts that aim at supporting industry-relevant research on self-adaptation. The guidelines that are grounded on data obtained from a survey with practitioners were derived during working sessions at the 17th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. Artifact providers can use the guidelines for aligning future artifacts with industry needs; they can also be used to evaluate the industrial relevance of existing artifacts. We also propose an artifact template.

1. INTRODUCTION

The term *artifact* has different meanings in the field of software engineering. In software development communities, artifacts refer to documents, deliverables, and work products [7]. Software engineering conferences increasingly use the term artifact in association with publications. Such artifacts can range from “functional” i.e., artifacts associated with the research that are documented and exercisable, and include appropriate evidence of verification and validation; to “replicated” i.e., the main results of the paper have been independently obtained in a subsequent study by a person or team other than the authors, without the use of author-supplied artifacts.¹ In this paper, we refer to an artifact as a tangible object that is created with the intention to drive research advances and compare and contrast alternative approaches.² Our particular focus is on artifacts that support research in the field of self-adaptation and its applicability in practice.

Generally, different types of artifacts can be provided. *Testbeds* refer to implementations of systems (or parts of systems) that support the study and evaluation of open challenges in self-adaptive systems. *Model problems* provide detailed specifications of sys-

tems that pose fundamental or characteristic challenges that self-adaptive systems should address, either in general or in specific domains. *Frameworks* refer to concrete or conceptual platforms with generic functionality that can be selectively specialized by implementing self-adaptation techniques or algorithms that are potentially useful in different contexts and use. *Datasets* (e.g., logging data, sensor data, system traces, survey raw data) can be used to develop, evaluate, and compare self-adaptation approaches. These common types of artifacts are representative classes of the existing artifact landscape.

Over the past decade, researchers in the field of self-adaptation developed about 30 artifacts [20]. A pioneering example is Znn.com [1] that provides a webserver system of a simplified news site. Znn.com’s testing environment simulates the slash-dot effect, which are periods of abnormally high traffic that overload the system. Another example artifact is TAS [23], which provides a service-based system, where services offered by third-party providers are dynamically composed into workflows delivering complex functionality under uncertainties such as changing workload and fluctuations in the network. DARTSim [15], instead, offers an artifact for the evaluation and comparison of adaptation approaches for smart cyber-physical systems that are subject to sensing errors among other uncertainties. Yet another example artifact is RoboMax [3], which provides an extensible repository of robotic mission adaptation exemplars. A recent study highlighted an increasing use of artifacts in the evaluation of research results, although the adoption by teams other than the developing team remains relatively low [2].

A recent analysis of the maturity of the field of self-adaptation [22] showed that we are currently in the phases of internal and external enhancement and exploration according to the model of Redwine and Riddle [16]. Hence, the industrial application of self-adaptation in real-life practical and commercial applications is of crucial importance to reach full maturity. Yet, it is currently not clear to what extent the existing artifacts support research activities of problems in self-adaptation that are relevant to industry.

To that end, this paper provides a set of guidelines for future artifact providers that aim at supporting industry-relevant research on self-adaptation. In addition, the guidelines enable assessing

¹<https://conf.researchr.org/track/icse-2022/icse-2022-artifact-evaluation>

²<https://conf.researchr.org/track/seams-2022/seams-2022-papers#Artifacts>

existing artifacts on their relevance to support industry-driven research. The guidelines are grounded on the data obtained from a large-scale survey with practitioners. The guidelines have been derived from the outcome of a series of working sessions during the in-person day at the 17th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2022.³ Note that alignment of research with industry has been topic of study in other domains, for instance in control [17] or in computing systems in general, see e.g., [4, 18].

The remainder of this paper is structured as follows. In Section 2, we present the process we followed to define the guidelines presented in this paper, and we summarize the data that was used as input for the working sessions. Section 3 then presents the guidelines. In Section 4, we present a template to describe industry-relevant artifacts on self-adaptation. Finally, we conclude the paper in Section 5.

2. BACKGROUND

In this section, we give an overview of the process we followed to create the guidelines for artifacts, together with a brief summary of the data used to define the guidelines.

2.1 Process used to Derive the Guidelines

Figure 1 shows the process we followed to derive the guidelines for artifacts.

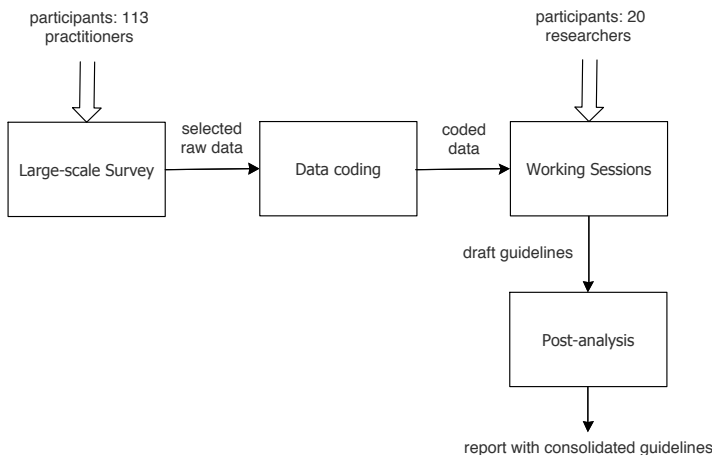


Figure 1: Process used to derive guidelines.

In a recent large-scale survey, we collected data from 113 practitioners⁴ about: (a) the drivers and motivations to apply self-adaptation in industrial software, (b) the use cases and approaches used to realize self-adaptation, (c) the problems and challenges practitioners face, and (d) the opportunities practitioners see for self-adaptation. For a summary of the initial results, we refer to [24]. When defining the guidelines, we have used raw data on three specific topics: (a) difficulties and challenges practitioners face when applying self-adaptation, (b) problems for which they would appreciate support from researchers, and (c) opportunities for applying self-adaptation. In addition, we used categories and codes that were obtained from the raw data through inductive

³<https://conf.researchr.org/track/seams-2022/seams-2022-papers>

⁴The participants range from different companies across the globe, active in a variety of domains; they have various roles in software engineering with a variety of experience (67% have at least 9 years of experience) [24].

reasoning (i.e., moving from specific fragments in comments to general concepts) [19].

The coded data (elaborated below) was then used as input for the working sessions that were organized during the in-person part of SEAMS 2022. Concretely, we divided the group of 20 participants in three breakout groups. Each group drafted a set of guidelines during two sessions of 1.5 hours. Finally, the insights were brought together and discussed during a plenary session of one hour. After the event, we consolidated the guidelines and documented the results in this paper.

2.2 Survey Data

We summarize the results of the coding of the survey data by highlighting (a) difficulties and challenges practitioners face when applying self-adaptation, (b) problems for which they would appreciate support from researchers, and (c) opportunities for applying self-adaptation.

Of all the surveyed practitioners, 35% regularly face difficulties and challenges with applying self-adaptation, while 40% report that they sometimes face problems. Table 1 shows the categories and codes collected from the answers. The table shows that practitioners face difficulties and challenges across different aspects of software engineering, from lifecycle issues to design, engineering processes, and runtime.

Table 1: Difficulties and challenges in self-adaptive systems as perceived by practitioners

Categories and codes	#
Lifecycle issues	22
Tuning/debugging	10
Limitations of tools/methods	7
System/environment evolution	5
Design challenges	20
Reliable/optimal design	10
Design complexity	9
Security	1
People and process issues	14
Skills/experience	8
Process and management	5
Documentation	1
Runtime challenges	13
Runtime uncertainty	6
Data collection/evaluation	5
Delayed/missing runtime changes	2

Twenty percent of the practitioners express that they would frequently appreciate support from researchers to tackle problems they face, while 23% would sometimes appreciate support. Table 2 shows the categories and codes collected from the answers. The table shows that practitioners would appreciate support for tackling a variety of problems, from engineering to providing guarantees of trust, user interaction-related solutions, management of data, and the realization of a number of advanced features.

Table 3 shows the opportunities practitioners see for applying self-adaptation that are currently not exploited. The main opportunities are related to system activities and system properties. Other opportunities concern engineering activities and human involvement in the self-adaptation process.

Table 2: Problems for which practitioners would appreciate support from researchers

Categories and codes	#
Engineering	22
Architecture & reuse	7
Adoption	7
Tools	4
Frameworks	4
Guarantees	19
Trustworthiness	9
Unknowns	5
Security	5
User interaction	18
User experience	7
Automation	7
User involvement	4
Data	18
Govern data	8
Assess data	6
Machine learning	4
Advanced features	13
Goals	6
Mechanisms	4
Testing	3

3. GUIDELINES

We now present the consolidated guidelines. Table 4 provides an overview of the guidelines that are divided in three groups: required, recommended, and desirable guidelines. We discuss each group now in detail.

3.1 Required Guidelines

This first group of guidelines are required for industry-relevant artifacts in self-adaptation. Yet, which of these guidelines are relevant and apply depends on the type of artifact and its concrete purpose.

P1: *Artifacts should take into account relevant surrounding elements*

Crucial for industry-relevant artifacts is that they take into account the context in which self-adaptation is applied. We refine this guideline in three concrete sub-guidelines.

P1.1: *Artifact should take into account management of real-world data*

There are different kinds of data that a self-adaptive system must manage during operation. Some are inherent to the functionality of the system, such as requests from users to a website and sensor readings in a cyber-physical system. Others are necessary for the self-adaptation process itself, such as monitoring data collected for analysis and planning. It is important that artifacts can manage data with characteristics representative of industrial use, including volume, rate, noise, etc. Otherwise, research evaluated with artifacts may not be able to deal with the data needs of industrial use.

P1.2: *Artifacts should facilitate humans-on-the-loop*

The initial research in engineering self-adaptive systems was mainly concerned with automating tasks [9, 12, 21]. In this research, the role of humans was in essence limited to providing

Table 3: Opportunities for self-adaptation in industry that are currently not exploited as mentioned by practitioners

Categories and codes	#
System activities	26
Autonomous operation	10
Autoscaling	8
AutoML & data management	8
System properties	20
Qualities	12
Cyber-security	8
Engineering activities	15
DevOps & maintenance	10
Adaptation patterns & libraries	5
Human involvement	5
Personalization	4
Human-machine interaction	1

high-level goals that the system should fulfill. Later, the role of humans-*in-the-loop* became prominent [5]. In this perspective, humans can take different roles, from providing input to supporting decision-making and guiding adaptation actions. Yet, several researchers pointed out the potential implications and risks of uncertainty caused by humans-*in-the-loop* [6, 11]. On the other hand, practitioners point to the need for humans-*on-the-loop* in self-adaptive systems. In industrial applications, humans should not be involved in the adaptation process all the time to avoid that they become a bottleneck. The preferred role of humans is instead supervision to achieve trustworthiness, avoid misuse or disuse of the automated system. In this regard, artifacts should consider different stakeholder roles: Users, operators, business owners, customers, etc., and emphasize their concerns such as feasibility, added value, and Return On Investment (ROI).

P1.3: *Artifacts should take into account industrial adoption of newly proposed solutions*

Artifacts are developed to evaluate and demonstrate new research results, ultimately with the goal of advancing the state-of-practice. It is however well known that industry will only adopt research solutions if they are demonstrated to be effective in the context of real operational requirements. As such, transferring research results to the field requires organizations to conduct a preliminary evaluation (e.g., in the form of an industrial prototype, testbed or case study), which has a cost that may be seen as a risk by managers and decision-makers. Artifacts can play a major role in the reduction of such costs, thereby facilitating the experimentation of new solutions by industry and promoting the adoption of research results. In practice, artifacts should be built in such a way that allow industry to easily develop proofs of concept to test the effectiveness of new research solutions in their specific settings and adopt the results in their practice.

P2: *Artifacts should take into account industry scale*

A representative scale is a key requirement for the industrial relevance of an artifact. Artifacts to support industry-relevant research should have a capacity and complexity that is characteristic of industrial practice. Taking into account real-world data as highlighted in guideline P1.1 is one specific facet of industry scale. Other aspects include the size of systems (e.g., a representative topology of an Internet of Things network), requirements of systems (e.g., representative requirements of health care systems in-

Table 4: Overview of guidelines for artifacts to support industry-relevant research on self-adaptation

Required Guidelines
P1: Artifacts should take into account relevant surrounding elements
P11: Artifact should take into account management of real-world data
P12: Artifacts should take into account humans-on-the-loop
P13: Artifacts should take into account industrial adoption of newly proposed solutions
P2: Artifacts should take into account industry scale
P3: Artifacts should provide industry-relevant metrics
P4: Artifacts should use an open policy
Recommended Guidelines
R1: Integration with standard technologies
R2: Process-related artifacts
R3: Generality and extensibility of artifacts
R4: Artifacts that facilitate collaborations across communities
Desirable Guidelines
D1: Artifacts may be built to exist in an ecosystem
D2: Artifacts may leverage open-source projects

cluding privacy, performance, and cost), the load exercised on systems (e.g., realistic workloads used by Cloud resources), and underlying infrastructure (e.g., representative messaging techniques for micro-service architectures).

P3: *Artifacts should provide industry-relevant metrics*

To quantify the benefits and costs of self-adaptation, artifacts should provide industry-relevant metrics. On the one hand, such metrics should address quality characteristics of software systems to measure the satisfaction of quality requirements. This allows practitioners and researchers to evaluate a self-adaptive system built with an artifact from a technical perspective. On the other hand, such metrics should also address a business-oriented perspective, for instance, to evaluate a self-adaptive system built with an artifact economically (e.g., added value and ROI as also suggested by guideline P1.2). This allows customers and business owners to assess self-adaptation solutions developed with a specific artifact. Moreover, practitioners and researchers may select suitable artifacts based on their technical and business interest by taking these metrics into account. This allows them to easily pick up connection points for leveraging the self-adaptation mechanisms.

P4: *Artifacts should use an open policy*

To facilitate adoption in industry, the artifacts need to (a) be openly available under clear licensing schemes, (b) leverage open standards (and preferably be based on standard technology, see guideline R1, and open-source projects, see guideline D2), and (c) ensure transparency. The transparency condition is crucial not only in terms of artifact trustworthiness, but also to support troubleshooting of defects in self-adaptation processes that are likely to be common in industry, given the reported shortage of expertise on the topic in the teams [24], and might require further process-related support (see guideline R2). The transparency also facilitates repeatability (i.e., same team, same experimental setup, for the sake of validation), reproducibility (i.e., different team, same experimental setup, for the sake of statistical significance), and replicability (i.e., different team, different experimental setup, for the sake of wider adoption). As artifacts may evolve over time, clear guidance on versioning and quality assurance is key to support industries willing to contribute to their extensions.

3.2 Recommended Guidelines

The second group of guidelines are recommended for industry-relevant artifacts in self-adaptation. We advise these guidelines as they will add to the relevance for industry.

R1: *Integration with standard technology*

Central to industrial relevance of artifacts is the use of standard technologies. This brings forward the following benefits:

- *Realism*: Artifacts developed with standard technologies allow for a closer level of realism due to the fact that the technologies used correspond to the actual technologies.⁵
- *Engagement*: Using standard technologies ensures that a larger share of people is familiar with how the artifact functions. Examples of standard technologies (at the time of writing) include: Docker, Real-time Operating System (ROS) for Robotics, Carla (self-driving simulator), etc.
- *Relevance*: The demonstrated benefits and evidence of techniques, solutions, or frameworks by using the artifact maps closer to industrial practice. More specifically, technology-specific settings of problems or opportunities put forward by industry connect more to the artifact when using the same technologies.

R2: *Process-related artifacts*

Process-related artifacts are those that are targeted at helping apply self-adaptation in industrial settings. Rather than tools that implement or support self-adaptation, these kinds of artifacts focus on software engineering processes (e.g., checklists) for designing, testing, and deploying them. For instance:

- *Design*: Guidance is needed on how to ensure that when applying self-adaptation in a particular context, that the

⁵Besides technologies used, real-world data traces, realistic uncertainty profiles, etc. should be used to ensure industry relevance.

system is complete and can provide a certain level of assurances. This requires that the right state is being modeled and monitored and the feedback loop is properly developed by combining tactics. E.g., it is not necessary to monitor elements that do not contribute to the decision-making about the system; similarly, if the self-adaptive system needs information in the planning or analysis process, this information needs to be stored in the knowledge that the self-adaptive system has about the system it is managing, and/or that appropriate monitoring is in place to retrieve that information in a timely and accurate way.

- *Testing*: In order to practically apply self-adaptation, we need strategies for testing self-adaptive systems aligned with current practice, definitions of standard test metrics like coverage in the context of self-adaptation, and guidelines on how to provide a testing environment that drives the adaptations to be tested. Some of the research behind this has been done in [8, 10, 13], but artifacts that can be used to operationalize self-adaptive testing would be helpful in guiding industry on how and what to apply.
- *Deployment*: Once an adaptation decision is made, the system needs to enact the adaptation. This calls for strategies to apply adaptation actions in a safe way. Artifacts to integrate the execution of system adaptation aligned with state of the practice infrastructure for continuous deployment are key in helping practitioners with closing the feedback loop of self-adaptation.

Process related artifacts may rely on, or take the form of best-practices documents, checklists for practitioners to use to ensure they deal with appropriate and common steps for self-adaptation, or test/deployment scripts that practitioners can use to manage CI/CD pipelines for self-adaptive systems.

R3: *Generality and extensibility of artifacts*

Generality and extensibility of artifacts add to their applicability and reusability. Therefore, it is important to define artifacts enabling their reuse both across application domains, such as robotics, automotive, and telecommunications, and/or axes of research/engineering activities, including requirements, architecture, verification, analysis, deployment, evolution, etc. Generality will leverage research results across domains and support the integration of research/engineering activities. Extensibility, on the other hand, enables the extension of existing artifacts to different domains and/or research/engineering activities. Note that this does not imply that domain-specific artifacts cannot be effective and useful. For instance, RoboMax [3], the artifact mentioned before, captures key sources of uncertainty and adaptation concerns of missions for the domain of robotic systems. Such domain-specific knowledge and know-how can add to the efficiency of applying an artifact.

R4: *Artifacts that facilitate collaborations across communities and with industry*

Industry requires simple and easily composable building blocks that can be shared for experimentation. The motivation for business gain could come from the composition of innovative ideas that have been developed by members of different scientific communities and mapped into robust workflows. Furthermore, science makes better progress when we have diversity of ideas and the people who generate the ideas [14]. To enable cross-community collaborations and integration with industry, researchers should make it easy for others to build on top

of the ideas and artifacts they generate. This calls for an infrastructure that easily facilitates automatic evaluation of methods/techniques/algorithms/etc. against benchmarks. Such an infrastructure will enable others to engage, while providing input and feedback from these diverse set of people. Several examples exist that have already demonstrated the added value; we give two examples:

- DLRM⁶ is a platform to which researchers from Computer Architecture and Systems and Recommendation Systems actively contribute. DLRM was developed by academics while placed as interns at Facebook.
- MLCommons⁷ supported the development of multiple important benchmarks such as MLPerf and TinyML for researchers interested in machine learning for systems and systems for machine learning.

In sum, we need an infrastructure, preferably open source, that makes it easier for researchers to share innovative ideas, for developers to experiment with ideas over a short time window (e.g., a weekend), and for software companies to make the ideas production-ready. This creates opportunities for more researchers, within and across different communities, to engage, and for industry to contribute to industry-relevant research efforts and engage with the research communities.

3.3 Desirable Guidelines

The third group of guidelines are desirable for industry-relevant artifacts in self-adaptation. These guidelines are worth having as they will add to the attractiveness for industry.

D1: *Artifacts may be built to exist in an ecosystem*

Desirable properties of artifacts for industry-relevant research on self-adaptation are extensibility and integratability, so that they can accommodate surrounding elements (compare also guideline R3). In the past, the focus of self-adaptive systems has been on mechanisms to adapt systems' behaviors or structures, whereas surrounding elements (e.g., data management, human-machine interaction, and integration with basic infrastructure) have received little attention (see guideline P1). To integrate surrounding elements, preferable in a lightweight way, it is desirable not to focus on specifics of self-adaptive systems in isolation, but to design artifacts to co-exist within an ecosystem. To support this guideline, for example, data management solutions may be provided as reusable components. Moreover, artifacts should be designed so that they can be extended and integrated with other systems (see also guideline R4). This will add to a holistic perspective, which is a key aspect for industrial relevance of artifacts.

D2: *Artifacts may leverage open-source projects*

To increase the industrial relevance of artifacts and their acceptance by industry, they would benefit from relying on well-known and robust projects. The use of open source projects may help artifacts to have a more stable underlying infrastructure. On the one hand, relying on existing infrastructure from open-source projects (e.g., containers, continuous integration tools, frameworks) reduces the learning curve for practitioners, and eases deployment and integration of the artifacts solutions with their technology

⁶<https://ai.facebook.com/blog/dlrm-an-advanced-open-source-deep-learning-recommendation-model/>

⁷<https://mlcommons.org/en/>

Table 5: Proposed template for industry-relevant artifacts to support research on self-adaptive systems

Element	Brief description
Problem	A brief description of the industry-relevant problem tackled by the artifact
Prerequisites	Requirements or constraints for using the artifact and the research results obtained with it
Technologies	Relevant state-of-the-practice technologies associated with the artifact
Outcome	The expected outcome of using the artifact and adopting the solution in practice (e.g., quantifiable by metrics)
Adoption	The relevant aspects required to adopt the research results obtained by using the artifact
Example	An exemplary solution to the problem and how it was adopted
Related artifacts	Artifacts that address similar problems or that can be combined to enrich the solution
Maintenance	A version history, and a reference to the maintainer of the artifact who keeps it up-to-date and responds to change requests

base. On the other hand, using existing projects as a base, can lead artifacts to reach a readiness level accepted by industry faster. Finally, we note that adopting an open-source process for the development of artifacts can be beneficial for their evolution and customization. As artifacts become more accepted, particular industries can customize a base artifact for their particular needs. As the artifacts evolve, new features can be integrated easily with branched and customized development.

4. TEMPLATE FOR INDUSTRY-RELEVANT ARTIFACTS

To support industry-relevant research on self-adaptation, artifact descriptions would benefit from a standard structure, similarly to the uniform descriptions of design patterns. Standard template-based descriptions support both researchers and practitioners in searching, identifying, understanding, and comparing artifacts that are relevant for a given problem. Otherwise, the artifacts themselves have to be investigated, analyzed, and even executed to assess their relevance for the given problem and technology, which is a difficult and time-consuming effort.

Table 5 shows the proposed template for industry-relevant artifacts.

The elements of the template were derived from the consolidated guidelines. The top part considers the basic elements to describe an artifact, from a description of the problem to the adoption of research results obtained with the artifact. The elements of the middle part provide an example solution and related artifacts that put the artifact in a broader context. Finally, the bottom part captures maintenance of the artifact in face of the rapid progress of technology in industry. Up-to-date maintained artifacts are essential to industrial relevance research. We anticipate that different types of templates might be needed to accommodate the description of different types of artifacts, such as testbeds, frameworks, datasets, and checklists. Finally, although the template elements proposed in Table 5 were derived based on data obtained from practitioners in the context of self-adaptation, the elements of the template are generic and applicable to industry-relevant research in other areas. Yet, more study is required to underpin this observation.

5. CONCLUSIONS

Artifacts are key to systematic evaluation of new research results in a field of interest. We provided a set of empirically grounded guidelines for artifacts that aim at supporting industry-relevant

research in the field of self-adaptation. The guidelines can be used by artifact providers for aligning future artifacts with industry needs as well as to assess the industrial relevance of existing artifacts. Finally, we proposed a template for describing industry-relevant artifacts to support research on self-adaptive systems. We hope that this paper will support the research community in self-adaptation to further enhance its maturity, the broader software engineering community to engage across their borders, and practitioners in adopting research results to build better software systems.

Acknowledgment

We are grateful to the practitioners that participated in the survey that we used as the basis for the presented guidelines.

6. REFERENCES

- [1] Znn.com. <https://github.com/cmu-able/znn>, 6/2022.
- [2] Shubham A. Analysing and assessing self-adaptive systems and designing a general criterion to evaluate sas artifacts. *MSc Thesis Vrije Universiteit Amsterdam*, 2020.
- [3] M. Askarpour, C. Tsigkanos, C. Menghi, R. Calinescu, P. Pelliccione, S. García, R. Caldas, T. von Oertzen, M. Wimmer, L. Berardinelli, M. Rossi, M. Bersani, and G. Rodrigues. Robomax: Robotic mission adaptation exemplars. In *International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2021.
- [4] N. Boules, K. Douglas, S. Feldman, L. Fix, G. Hager, B. Hailpern, M. Hebert, D. Lopresti, B. Mynatt, C. Rossbach, and H. Wright. The future of computing research: Industry-academic collaborations. *Computing Community Consortium*, Version 2:1–19, 6/2022.
- [5] J. Cámara, G. Moreno, and D. Garlan. Reasoning about human participation in self-adaptive systems. In *IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2015.
- [6] N. Esfahani and S. Malek. Uncertainty in self-adaptive software systems. In R. de Lemos, H. Giese, H. Müller, et al., editors, *Software Engineering for Self-Adaptive Systems II: International Seminar, Dagstuhl Castle, Germany, October 24-29, 2010 Revised Selected and Invited Papers*, pages 214–238. Springer Berlin Heidelberg, 2013.
- [7] D. Fernández, W. Böhm, A. Vogelsang, J. Mund, M. Broy, M. Kuhrmann, and T. Weyer. Artefacts in software engineering: a fundamental positioning. *Softw. Syst. Model.*, 18(5):2777–2786, 2019.

- [8] E. Fredericks and B. Cheng. Automated generation of adaptive test plans for self-adaptive systems. In *IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2015.
- [9] D. Garlan, S.W. Cheng, A.C. Huang, et al. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, 37(10):46–54, 2004.
- [10] O. Gheibi and D. Weyns. Lifelong self-adaptation: Self-adaptation meets lifelong machine learning. In *IEEE/ACM 17th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2022.
- [11] S. Hezavehi, D. Weyns, P. Avgeriou, R. Calinescu, R. Mirandola, and D. Perez-Palacin. Uncertainty in self-adaptive systems: A research community perspective. *ACM Transactions on Autonomous and Adaptive Systems*, 15(4), 2021.
- [12] J. Kephart and D. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
- [13] M. Langford and B. Cheng. Know what you know: Predicting behavior for learning-enabled systems when facing uncertainty. In *International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2021.
- [14] D. Medin and C. Lee. Diversity makes better science. *APS Observer*, 25(5), 2012.
- [15] G. Moreno, C. Kinneer, A. Pandey, and D. Garlan. Dartsim: An exemplar for evaluation and comparison of self-adaptation approaches for smart cyber-physical systems. In *IEEE/ACM 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2019.
- [16] S. Redwine and W. Riddle. Software technology maturation. In *8th International Conference on Software Engineering*, page 189–200, Washington, DC, USA, 1985. IEEE.
- [17] T. Samad, M. Bauer, S. Bortoff, S. Di Cairano, L. Fagiano, P. Odgaard, R. Rhinehart, R. Sánchez-Peña, A. Serbezov, F. Ankersen, P. Goupil, B. Grosman, M. Heertjes, I. Mareels, and R. Sosseh. Industry engagement with control research: Perspective and messages. *Annual Reviews in Control*, 49:1–14, 2020.
- [18] A. Scandura and S. Iammarino. Academic engagement with industry: the role of research quality and experience. *Technology Transfer*, Preprint:1–37, 2021.
- [19] K. Stol, P. Ralph, and B. Fitzgerald. Grounded theory in software engineering research: A critical review and guidelines. In *38th International conference on Software Engineering (ICSE)*, pages 120–131, 2016.
- [20] T. Vogel. Catalog of SEAMS artifacts: <http://self-adaptive.org/exemplars>. 6/2022.
- [21] D. Weyns. Software engineering of self-adaptive systems. In Sungdeok Cha, Richard N. Taylor, and Kyochul Kang, editors, *Handbook of Software Engineering*, pages 399–443. Springer International Publishing, Cham, 2019.
- [22] D. Weyns. *An Introduction to Self-adaptive Systems: A Contemporary Software Engineering Perspective*. John Wiley & Sons, 2021.
- [23] D. Weyns and R. Calinescu. Tele assistance: A self-adaptive service-based system exemplar. In *IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2015.
- [24] D. Weyns, I. Gerostathopoulos, N. Abbas, J. Andersson, S. Biffl, P. Brada, T. Bures, A. Di Salle, P. Lago, A. Musil, J. Musil, and P. Pelliccione. Preliminary results of a survey on the use of self-adaptation in industry. In *International Symposium on Software Engineering of Adaptive and Self-Managing Systems*. <https://arxiv.org/abs/2204.06816>, 2022.