ISSN 2186-7437

NII Shonan Meeting Report

No. 241

Patterns and Practices in AI Engineering and Governance

Hironori Washizaki
Foutse Khomh
Qinghua Lu
Shaukat Ali

February 9–12, 2026

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Patterns and Practices in AI Engineering and Governance

Organizers:
Hironori Washizaki (Waseda University and
National Institute of Informatics, Japan)
Foutse Khomh (Polytechnique Montreal / MILA, Canada)
Qinghua Lu (DATA61, CSIRO, Australia)
Shaukat Ali (Simula Research Laboratory, Oslo, Norway)

February 9–12, 2026

Context and goal

The popularity of artificial intelligence (AI), including machine learning (ML) techniques, has increased in recent years. AI is used in many domains, including cybersecurity, the Internet of Things (IoT), digital twins and autonomous cars, and is expanding its impact in scientific research, consumer assistants, and enterprise services through advancements in Generative AI (GenAI). Many works have investigated the mathematics and algorithms on which the AI techniques and models are built, but few have examined system engineering as well as their governance, which ensures AI systems are built, used, and managed to maximize benefits and prevent harms. AI engineering and governance needs to bring together diverse stakeholders across AI algorithms, data science, software/system engineering, compliance, legal, and business teams.

In AI software engineering and governance, there is often a gap between high-level abstract principles and low-level concrete tools and rules. Patterns encapsulating recurrent problems and corresponding solutions under particular contexts and pattern languages as organized and coherent patterns can fill such gaps, resulting in a common "language" for various stakeholders involved in often interdisciplinary AI software systems development and governance. Researchers and practitioners study best practices for engineering and governing AI/ML systems to address issues in AI and ML techniques as well as processes, policies, and tools for trustworthy, responsible and safe AI system development and management. Such practices are often formalized as patterns and pattern languages. Major examples are:

- AI architecture and design patterns, such as software engineering patterns for ML applications [7], ML design patterns [2] and agent design patterns [3].
- AI assurance argument patterns, such as safety case patterns for ML systems [9] and security argument patterns for DNN [10]

- Responsible AI engineering and governance patterns, such as patterns for creating trustworthy and safe AI systems [4]
- AI development and management practices, such as lifecycle phase practices [5]
- Prompt engineering patterns such as prompt pattern catalogue and taxonomy [8, 6]
- RAG (Retrieval Augmented Generation) patterns [1]

While AI engineering and governance patterns have been documented, there's still much to uncover in this landscape. This limited understanding hampers adoption, preventing the realization of their full potential.

The goal of the meeting was to bring together software engineering and AI experts from academia and industry, with a special focus on the theoretical, social, technological, and practical advances and issues related to patterns and practices in AI engineering and governance.

Details of the meeting

The meeting had two types of sessions: The first type of session had short presentations in which each participant introduced their research work and listed ideas, challenges, and research directions they wished to discuss during the meeting. Such presentations were scheduled during the first two days of the meeting. The second type of session consisted of intensive discussions among sub-groups of participants. The topics of the discussions were decided at the meeting from those proposed by participants during their short talks; however, the organizers prepared some possible discussion topics beforehand. The discussions proceeded in two phases: first exploring different topics, and then deepening the discussion on a restricted set of selected topics.

- **Exploration phase:** In the first instances of this type of session, the meeting used a dedicated method to guide the discussion. The plan was to use a variation of the world café method[2]: this allowed participants to move across the different sub-groups and thus get to know the different initial topics. In addition to general discussions on the grand challenges and future directions, potential topics also included pattern-specific ones, such as associating existing patterns to organize a pattern system beyond individual areas, mapping patterns and practices onto AI engineering and governance activities, and mining new patterns and practices. Afterwards, each participant decided which topic they would like to explore. At this stage, some topics that did not receive much interest were dropped.
- **Deepening phase:** The remaining sessions were dedicated to intensive discussion on the selected topics. Ultimately, each group came up with a research agenda for its topic. Afterwards, we planned to develop a research agenda for patterns and practices in AI engineering and governance, and to prepare a book proposal based on the agenda that was discussed and

[1] https://humanloop.com/blog/rag-architectures
[2] http://www.theworldcafe.com/method.html

agreed upon by the participants. We plan to apply to the “Call for Book Proposals” provided by the Shonan meeting organisation.

Overview of the meeting

The meeting began with an introduction by the organizers that outlined the overall scope, schedule, potential outcomes, and other practical matters. This was followed by a brief introduction to each participant—approximately 2 minutes each, with or without slides to present themselves, outline the topics they are working on, and share their expectations for the seminar. More detailed presentations from selected participants followed this, each lasting approximately 15 minutes. The titles and abstracts of these talks are presented in the "Overview of Talks" section. The result of these talks was the identification of key topics to discuss, as well as the formation of groups for focused discussion for the rest of the seminar. These topics were identified for discussions:

(i) Evolution and Operations: The Strategic Future of AI Engineering

(ii) Pattern Evaluation in AI Engineering

(iii) Human AI Interaction

(iv) Pattern-Driven Co-Design with Human and AI Agents: A Lifecycle Approach

In the rest of the report, we first present the meeting schedule, then the abstracts of the talks, and finally a summary of the group discussions.

Meeting Schedule

February 8, 2026, Sunday

- 19:00- Welcome reception

February 9, 2026, Monday

- 9:00-9:15 Opening
- 9:15-10:30 Introduction. 2 mins for each participant
- 10:30-11:00 Group Photo Shooting and break
- 11:00-12:00 Presentations & Discussion
- 12:00-13:30 Lunch
- 13:30-15:00 Presentations & Discussion
- 15:00-15:30 Break
- 15:30-17:00 Presentations & Discussion
- 17:00-17:15 Wrap-up of first day
- 18:00- Dinner

February 10, 2026, Tuesday

- 9:00-10:30 Discussion on an initial set of topics for the discussion in the sub-groups
- 10:30-11:00 Break
- 11:00-12:00 Deciding the initial set of topics with preference voting, followed by wrap-up of second day
- 12:00-13:30 Lunch
- 13:30- Excursion (including Dinner)

February 11, 2026, Wednesday

- 9:00-10:30 Discussion of the initial selected topics in sub-groups (exploration phase)
- 10:30-11:00 Break
- 11:00-11:45 Discussion of the initial selected topics in sub-groups (exploration phase)
- 11:45-12:00 Summary of the discussions of the different sub-groups and selection of a restricted set of topics to be discussed in the deepening phase. Each participant is allocated to a specific topic
- 12:00-13:30 Lunch

- 13:30-15:00 Discussion of the selected topics in sub-groups (deepening phase)
- 15:00-15:30 Break
- 15:30-16:45 Discussion of the selected topics in sub-groups (deepening phase)
- 16:45-17:00 Wrap-up of third day
- 18:00- Dinner
- Night session

February 12, 2026, Thursday

- 9:00-10:30 Discussion of the selected topics in sub-groups (deepening phase)
- 10:30-11:00 Break
- 11:00-12:00 Presentations of the work of the individual sub-groups. Final planning of future work and wrap-up
- 12:00-13:00 Lunch

Overview of Talks

Architectural Patterns for Federated Learning

Catia Trubiani, Gran Sasso Science Institute, Italy

Federated Learning has emerged as a promising paradigm that enables collaborative model training while preserving data privacy. However, the design of Federated Learning systems necessitates non-trivial architectural choices to address several challenges, including system efficiency and learning accuracy. The objective of this talk is to discuss how to empower software architects in their task of evaluating the design of FL systems while deciding which architectural alternatives are more beneficial in their context of adoption.

AI Agent Construction Patterns: From the Perspective of Personas and AI Clones

Naoyasu Ubayashi, Waseda University, Japan

With the rapid advancement of generative AI, collaborative work between humans and AI systems has become increasingly common. AI capabilities are often realized through AI agents; however, the methodological foundations for constructing such agents remain insufficiently developed. Personas and AI clones offer a promising conceptual and practical basis for addressing this gap. In particular, personas have long been used in software engineering as a technique for eliciting user requirements. The research community has proposed a wide range of methods for generating personas and AI clones, including data-driven approaches that leverage demographic information and behavioral logs, as well as heuristic and expert-guided approaches. Despite their diversity, these methods exhibit recurring structural and procedural characteristics. To systematize this accumulated knowledge and support future AI-assisted software development, we propose organizing these recurring elements into a structured catalogue of AI agent construction patterns. Furthermore, as interactions between humans and AI systems increasingly raise conceptual and ethical concerns, the construction of AI agents should be informed by philosophical considerations.

Architectural Patterns for Human–AI Interaction in Socio-Technical Systems: Education as a Lens

Sherry Xu, CSIRO, Australia

As generative AI systems increase in complexity and are embedded more deeply into real-world workflows, human–AI interaction becomes an increasingly salient architectural concern. Architectural decisions shape how control, responsibility, and context are distributed between humans and AI over time, influencing software quality attributes such as scalability, auditability, adaptability, and maintainability under multiple constraints.

This working session brings participants together to explore and refine architectural patterns for human–AI interaction in socio-technical systems, using education as a constraint-rich domain. Education exposes technical (e.g., scalability, consistency), regulatory (e.g., privacy, compliance), institutional (e.g., internal procedure), and normative (e.g., interpretive authority, pedagogical validity) constraints that act as architectural drivers. While grounded in education, the patterns discussed are intended to generalise to other AI-enabled systems operating under strong organisational and societal constraints.

Architecturally Significant AI Patterns: industrial studies and beyond

Henry Muccini, University of L'Aquila, Italy

This presentation explores the landscape of architecturally significant patterns within agentic systems driven by Large Language Models (LLMs). While numerous theoretical design patterns—such as ReAct, Reflexion, and Toolformer—have emerged, their industrial applicability and empirical impact remain under-investigated. Our research addresses these gaps through a multidimensional study:

- In a first study, we made a qualitative analysis of 138 practitioner presentations identifies a taxonomy of recurring industrial patterns categorized into coordination, control, reliability, and evolution.
- On a different ongoing study, we are utilizing repository mining across approximately 13,000 software projects to quantify the frequency, combination, and non-compliance of established agentic patterns in production code.
- Furthermore, we are working to introduce an empirical framework to evaluate the energy-consumption profiles and task performance of varied LLM planning architectures, holding operational variables constant to isolate architectural effects.

By investigating these open questions, the research aims to provide a systematic methodology for architects to implement energy-efficient and reliable multi-agent systems while advancing the state of Green Software engineering in the era of autonomous AI.

AI Patterns Engineering

Hironori Washizaki, Waseda University, Japan

In AI and machine learning (ML) software engineering, there is often a gap between high-level abstract concepts and principles and low-level concrete tools and cases. Patterns encapsulating recurrent problems and corresponding solutions within particular contexts and pattern languages, organized and coherent, can fill such gaps, resulting in a common "language" for various stakeholders involved in the often interdisciplinary development of AI/ML software systems. This talk first introduces concepts of pattern languages in general and major AI/ML software engineering patterns, including ML architecture and design patterns, ML assurance argument patterns, and LLM/VLM multi-agent design patterns. Then, the talk explains emerging engineering approaches for extracting, detecting, and applying AI/ML patterns with the support of LLMs.

Better Together: How Humans and AI Can Co-Create Software Designs

Rick Kazman, University of Hawaii, USA

In this talk, I will argue that we are entering a golden age of software design–one that capitalizes on the unique capabilities of both AI and humans. We have seen that creativity in software design benefits from a division of labor: AI excels at divergent thinking–rapidly generating ideas and alternatives—while humans excel at convergent thinking–evaluating, refining, and selecting solutions while considering context and constraints. The result is not a replacement of human creativity in design, but rather an amplification of it.

Agentic AI for Software: from Scale to Trust

Abhik Roychoudhury, National University of Singapore, Singapore

In this talk, I will discuss how the use of AI agents in software engineering is turning the focus from building systems at scale to building systems with trust. This defines a new set of problems to look into, and using the agents in innovative ways to engender trust.

Patterns, Anti-Patterns vs. Dark Patterns in Agentic AI

Cesare Pautasso, University of Lugano, Switzerland

The rapid integration of autonomous agents into everyday software has amplified the importance of design patterns that guide trustworthy behavior, while simultaneously exposing users to subtle manipulations that exploit their trust. In this talk we attempt to classify and compare different Agentic AI practices and discuss where to draw the line in terms of their ethical implications.

Designing Agentic Systems: A Domain-Driven Approach to Language-Centric AI

Joseph Yoder, The Refactory, USA

Large Language Models (LLMs) have transformed computing by making natural language the primary interface between humans and machines. This shift has sparked widespread experimentation, yet the design of agentic systems remains ad hoc and fragile. In this talk, we propose a principled foundation for agent design by adapting key constructs from Domain-Driven Design (DDD). We reinterpret Bounded Context to define the scope of agent responsibility, and Ubiquitous Language as the vocabulary agents use to reason, prompt, and interact with users, other agents, and external APIs. By embedding agentic design within established system architecture practices, we offer a repeatable and rigorous methodology for building hybrid conventional and AI systems—one that treats language not just as an interface, but as infrastructure.

Towards AI-driven organizations

Jan Bosch, University of Groningen & Chalmers University of Technology, Netherlands, and Helena Holmström Olsson, Malmö University, Sweden

There are few technologies, if any, that have the potential to change the software-intensive industry to the extent that artificial intelligence (AI) is currently doing. Across industries, companies are adopting these technologies to improve productivity, to increase efficiency and to automate tasks. In products, AI is used for optimization and mass-customization. However, there are few examples of companies that use AI to reinvent and fundamentally change their existing practices. In this paper, we present results from an expert interview study in which we explore how AI is affecting the ways in which companies operate and what steps companies evolve through when advancing their use of AI in their development processes, in their products and in their business processes. The contribution of the paper is two-fold. First, we present the interview results that reflect the adoption and use of AI technologies. As part of our interviews, we also identify a set of key challenges that companies experience. Second, we present an inductively derived three-pronged maturity model that describes how companies transition from traditional towards AI-driven organizations.

A Pattern Language for Enterprise AI UX Design

Jung-Sing Jwo, Tunghai University, Taiwan and Yu Chin Cheng, National Taipei University of Technology, Taiwan

As enterprise AI spreads, human–AI interaction shifts from one-off tool use to long-term working relationships embedded in workflows. Usability alone is insufficient. HITL (Human-in-the-Loop) systems must satisfy both adoption and governance. Employees must use AI confidently, willingly, correctly (avoiding over-trust and undue distrust), and sustainably (governable, improvable, scalable). This talk proposes a pattern language for enterprise AI UX that formalizes these requirements into reusable design patterns. We define three core dimensions: predictability, controllability, and accountability, and three extensions: trust calibration, legitimacy, and non-adversarial governance, to reduce adoption friction, support workflow integration, and avoid shifting risk to frontline staff. This pattern language reframes enterprise AI UX from an interface-level usability problem into an organizationally acceptable, auditable, and continuously improvable framework for designing and governing human–AI working relationships.

Project Practice-Based Anti-patterns for AI Projects and Their Effective Utilization

Hironori Takeuchi, Musashi University, Japan

In this talk, we consider AI projects in which service systems are developed using machine learning (ML) techniques. From project practices, we have identified warning signs in AI projects, known as "bad smells," and developed anti-patterns based on them. As a result, we identified 18 bad smells, referenced

existing best practices as solutions, and developed 14 anti-patterns for ML projects. Based on the results of a survey on the identified bad smells, we found that some of the developed anti-patterns are related to specific practitioner types or project types. We will discuss the possibility of meta-knowledge for utilizing anti-patterns in AI project practices.

Challenges on Dynamic Documentation for LLM-Driven Development: Knowledge Structuring and Scalable Evaluation

Takashi Kobayashi, Institute of Science Tokyo, Japan

As LLMs increasingly generate code, developers shift toward reviewing and orchestrating tasks; making program comprehension the new bottleneck. Local documentation is a solved problem, but global context remains elusive. Dynamic documentation must be built on structured project knowledge, with RAG acting as a knowledge API rather than a search engine. Evaluating generated documents is not just a research concern—it's essential for reliable engineering. LLM-as-a-Judge is emerging as a scalable solution and is already gaining traction in SE research. To move forward, we must establish shared taxonomies and quality factors for dynamic documentation in LLM-era workflows.

Contextualizing AI agents with software and system models

Judith Michael, University of Regensburg, Germany

Agentic AI approaches for system engineering and operation promise flexible and adaptive behavior, but they often lack explicit structure, traceability, and alignment with system goals. This talk argues that software and systems models (such as UML, SysML, or DSLs capturing Personas or goals) provide essential context for AI agents by making assumptions, constraints, and domain semantics explicit. We present some insights on combining inductive AI components with deductive models and genAI with model-driven engineering. Examples from model-driven engineering, digital twins in education, and human–CPS interaction illustrate how genAI can both benefit from and support modeling tasks. Overall, contextualizing AI agents with software and systems models is key to building understandable, and trustworthy AI-enabled systems. However, patterns to support these practices still need to be identified.

Focus group on “Evolution and Operations: The Strategic Future of AI Engineering”

Participants: Henry Muccini, Abhik Roychoudhury, Rick Kazman, Fuyuki Ishikawa, Catia Trubiani, Helena Holmström Olsson, Jan Bosch, Shaukat Ali, Hironori Washizaki

The landscape of software engineering is undergoing a significant shift due the massive introduction of AI as a major innovator. The “Evolution and Operations” group at the Shonan meeting investigated the role of Patterns in the context of quality-driven design. More specifically, how patterns can be correctly engineered and evaluated, as well as how systems (including patterns) can be properly engineered.

To better understand pros and cons of AI in the context of Software and Patterns Engineering, the group ran a comprehensive SWOT (strengths, weaknesses, opportunities, and threats) analysis. The analysis revealed the transformative potential and inherent risks of GenAI and Agentic AI engineering:

- Strengths: AI can manage complexity beyond human capabilities and promises to provide orders-of-magnitude improvements in productivity. It also democratizes development by allowing non-experts to create valuable, customized systems and speeds up innovation through the rapid generation of MVPs.
- Weaknesses: Challenges include training bias, a lack of clear boundaries regarding what the technology can and cannot do, and the difficulty of understanding the ”why” behind generated content. Furthermore, a few leading tech companies currently control the capabilities roadmap.
- Opportunities: There is significant potential for runtime system generation and the exploration of support for quality attributes like robustness and performance. AI also offers new ways to handle regulatory compliance in safety-critical systems and interfaces with physical hardware (mechanics and electronics).
- Threats: Key concerns involve security risks, the potential for ”dark patterns,” and a heavy dependency on foundation model providers. There are also broader societal risks, such as the de-skilling of the workforce and sustainability issues regarding energy and knowledge retention.

Conclusion and Future Outlook: The strategic future of AI engineering depends on balancing the speed of innovation with rigorous quality assurance. To have a lasting impact on industry and society, the field must develop methods to evaluate both individual patterns and their broader systemic effects. Future efforts should focus on combining generative power with human knowledge to ensure safety and reliability.

Focus group on "Pattern Evaluation in AI Engineering"

Participants: Catia Trubiani, Hironori Washizaki

Patterns, in a sense of pattern languages, are basically extracted from the past, by abstracting and generalizing successful cases and practices, resulting in identifying recurrent problems and corresponding solutions under particular contexts. Due to the nature of patterns, the key question for managing patterns is how to evaluate patterns, which would be useful to be applicable to a future context with good impacts. This question is particularly crucial for GenAI engineering patterns if GenAI is used to generate solutions tailored to requirements and contexts at each time, rather than reusing past solutions, resulting in their limited proliferation.

The group outlines a comprehensive approach to evaluating patterns, especially AI patterns, addressing major characteristics of patterns through various methods. The characteristics include positive and negative impacts on quality attributes, such as AI utility and functionality (e.g., prediction accuracy) as well as AI system's performance and security, general applicability, and margin of solution.

Quality attributes: Existing metrics approaches, such as ISO/IEC 25010 quality model and the Goal-Question-Metric (GQM) method, can be applied through instantiations and emulations of AI systems with proper application of AI patterns.

General applicability: Through such emulation, the general applicability is also expected to be examined by a coverage of the potential instantiation space of the target AI pattern with proper sampling strategies.

Margin of solution: A tradeoff of the general applicability against the margin of solution should be important for AI patterns, which is expected to be examined by analyzing the number of available instantiation options (i.e., degree of freedom) and checking (mis)match and fitness between the pattern's context and specification and applicable targets' ones.

Available evaluation methods include the following:

- Experimental methods include emulation and metrics
- Observational methods include case studies and users' perceptions
- Analytical methods include argumentation and formal analysis
- Social validation includes a result of pattern adoption and experts' peer review

Focus group on "Human AI Interaction"

Participants: Sherry Xu, Marija Slavkovik, Xiaoning Du, Cesare Pautasso, Takashi Kobayashi, Shinpei Hayashi, Naoyasu Ubayashi, Foutse Khomh, Shaukat Ali, Catia Trubiani

The discussion group focused on human-AI interaction patterns. The discussion resulted in a framework that conceptualizes human–AI interaction as the interplay between three core elements: human characteristics, AI characteristics, and human-AI interaction structure. On the human side, behavior and decision-making are shaped by three key factors: domain expertise (high or low), trust level (high—for example, when under time pressure—or low, such as in a skeptical stance), and AI/IT expertise (high or low). These dimensions influence how individuals interpret, rely on, and supervise AI outputs. Depending on the context, humans may assume different roles—consumer, developer, operator, or regulator—each associated with distinct responsibilities and expectations within the interaction. On the AI side, systems are characterized by agency (ranging from low-level models to high-agency autonomous agents), connectivity (local deployment or commercial API-based access), inspectability (from glass-box transparency to black-box opacity), confidence level (if reported, it can be high or low, often evaluated relative to human expectations and trust), and dependability (high or low). AI systems primarily contribute through generation and reasoning capabilities. The interaction between human and AI is further structured by the level of automation—conceptualized across ten levels—which determines how authority and control are distributed between the two elements in practice. The group also evaluated the framework with several real use cases including conference paper review process, medical diagnosis, elderly care, educational bot, customer service, and openclaw.

Focus group on "Pattern-Driven Co-Design with Human and AI Agents: A Lifecycle Approach"

Participants: Hironori Takeuchi, Judith Michael, Yu Chin Cheng, Jung-Sing Jwo, and Joseph Yoder

This report synthesizes the collaborative work of group members focused on the topics knowledge and representation. We discussed how knowledge about patterns can be represented and shared effectively within the context of human-AI co-design processes. As AI agents become increasingly capable collaborators in software design, teams face a new challenge: how to create shared representations that allow humans and AI to co-design effectively. Traditional design patterns—rich in narrative, tacit knowledge, and conceptual framing—are optimized for human understanding but poorly suited for AI reasoning. GenAI and Agentic AI approaches need well-structured information about commonly used and known patterns to provide comprehensible and explainable AI-assisted design suggestions. Patterns could evolve, need to be extended, or have to be constrained over time based on experiences from application practice. This leads to the need for managing the **pattern lifecycle**. The pattern lifecycle describes how patterns evolve within a human–AI co-design environment. It begins with existing patterns and proceeds through assimilation, dual representation, co-creation, and co-enhancement. Each stage contributes to transforming patterns into shared cognitive artifacts that support collaborative design.

This **pattern repository** forms an important knowledge source for agentic GenAI to generate design artifacts. With a validated set of patterns in place, the lifecycle shifts from preparing knowledge to organizing and operationalizing it. The next stage involves selecting, adapting, and instantiating these patterns to form a coherent **Project Language** that reflects the requirements, constraints, and domain semantics of the system under development. This Project Language becomes the practical design vocabulary through which humans and AI agents jointly generate, evaluate, and refine system-level solutions for the specific project you are designing.

Figure 1 shows a "System Co-Creation Process," which provides a view of how humans and AI agents interact using common patterns and project-specific knowledge during the design of a software system. After GenAI has pre-learned pattern knowledge (step A), the co-creation process starts with (B.1) the specification of relevant domain knowledge, the project scope, initial requirements specifications, and optionally allowed patterns in this project context by the project team. This forms the initial context for GenAI agents (B.2) to create the first design artifacts (B.3). The project team evaluates these created artifacts (B.4) and provides a set of additional context information (B.5) for the next iteration loop on these artifacts by GenAI agents. In an iterative development process, these steps (B.2-B.5) are repeated to create a satisfactory system. While doing this, both the GenAI agents (C.1) and the project team (C.2) can create new patterns or adapt existing patterns that update the pattern repository. The project team can use a GenAI agent to assist with project-specific knowledge (D).

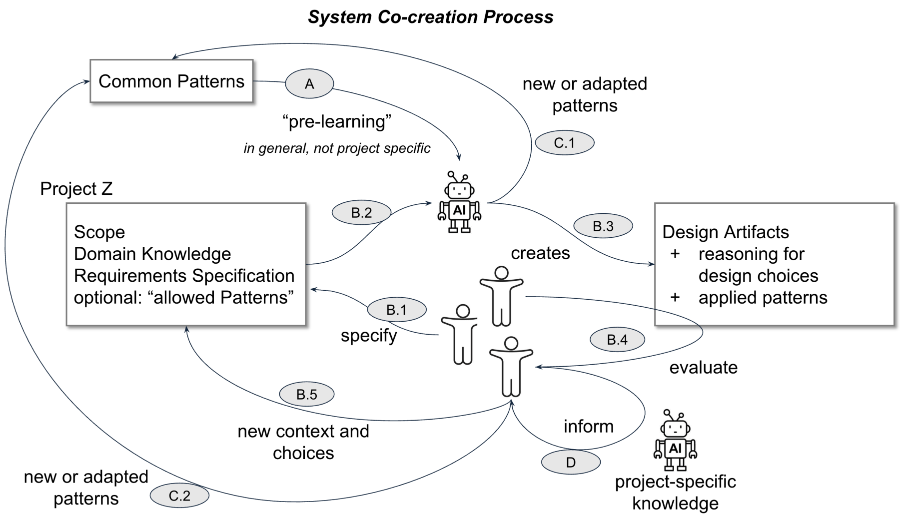

Figure 1: The envisioned system co-creation process

Each of these phases needs further research. For example, on which representations and knowledge about patterns are helpful for humans or GenAI agents, and on how to do the pre-learning. Additionally, it is important to validate the patterns and their correct application, on which agents are needed to create which parts of the artifacts, and how to validate the usefulness of AI-assistance in co-creation processes. Additionally, there needs to be concrete case studies that show the feasibility of such an approach.

Community Building and Expansion

This Shonan meeting provided an excellent opportunity to establish a community of researchers and practitioners from diverse areas, such as software and systems engineering, AI engineering, AI governance, and pattern languages, focused on patterns and practices for AI engineering and governance. To continue, strengthen, and expand this community, the meeting organizers and participants plan to hold related community events, including AI Patterns Tokyo 2026 [1], a seminar on AI and patterns held back-to-back with this Shonan Meeting, as well as other workshops and meetings, possibly co-located with major conferences or organized independently.

The organizers and participants also plan to further develop the research agendas initiated during the meeting, pursue joint research collaborations, prepare a book proposal for the Shonan Meeting Series, and explore other related opportunities, such as journal special issues and visionary papers.

Conclusion and Future Work

This Shonan meeting provided a collaborative forum for researchers and practitioners to exchange ideas and explore emerging challenges in patterns and practices for AI engineering and governance. Through presentations and focused group discussions, participants identified key research directions in AI pattern evaluation, human–AI interaction, pattern lifecycle management, and the strategic evolution of AI engineering. The discussions highlighted the need for systematic methods to represent, evaluate, and operationalize patterns in increasingly complex AI-driven systems.

As future work, the participants plan to further develop the research agendas initiated during the meeting, pursue joint research collaborations, prepare a book proposal for the Shonan Meeting Series, and explore additional dissemination opportunities such as journal special issues, visionary papers, and community events, including workshops, seminars, and future Shonan-style meetings.

References

[1] AI Patterns Tokyo 2026: Agentic & Responsible AI and Patterns Seminar, 2026. `https://smartse.connpass.com/event/382391/` [Last access: March 12th, 2026].

[2] V. Lakshmanan, S. Robinson, and M. Munn. *Machine Learning Design Patterns: Solutions to Common Challenges in Data Preparation, Model Building, and MLOps.* O'Reilly Media, Sebastopol, CA, 2020. First edition.

[3] Y. Liu, S. K. Lo, Q. Lu, L. Zhu, D. Zhao, X. Xu, S. Harrer, and J. Whittle. Agent design pattern catalogue: A collection of architectural patterns for foundation model based agents. *Journal of Systems and Software*, 220:112278, 2025.

[4] Q. Lu, L. Zhu, J. Whittle, and X. Xu. *Responsible AI: Best Practices for Creating Trustworthy AI Systems.* Addison-Wesley Professional, 2023.

[5] M. S. Rahman, F. Khomh, A. Hamidi, J. Cheng, G. Antoniol, and H. Washizaki. Machine learning application development: Practitioners' insights, 12 2021.

[6] Y. Sasaki, H. Washizaki, J. Li, N. Yoshioka, N. Ubayashi, and Y. Fukazawa. Landscape and taxonomy of prompt engineering patterns in software engineering. *IT Professional*, 27(1):41–49, 2025.

[7] H. Washizaki, F. Khomh, Y. Gueheneuc, H. Takeuchi, N. Natori, T. Doi, and S. Okuda. Software-engineering design patterns for machine learning applications, Mar. 2022.

[8] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, and D. C. Schmidt. A prompt pattern catalog to enhance prompt engineering with chatgpt, 2023.

[9] E. Wozniak, C. Carlan, E. Acar Celik, and H. Putzer. *A Safety Case Pattern for Systems with Machine Learning Components*, pages 370–382. 09 2020.

[10] M. Zeroual, B. Hamid, M. Adedjouma, and J. Jaskolka. Security argument patterns for deep neural network development. In *Proceedings of the 30th Conference on Pattern Languages of Programs (PLoP)*, pages 1–18, 2023.

List of Participants

- Shaukat Ali (Simula Research Laboratory, Norway)
- Jan Bosch (University of Groningen & Chalmers University of Technology, Netherlands)
- Yu-Chin Cheng (National Taipei University of Technology, Taiwan)
- Xiaoning Du (Monash University, Australia)
- Shinpei Hayashi (Institute of Science Tokyo, Japan)
- Helena Holmström Olsson (Malmö University, Sweden)
- Fuyuki Ishikawa (NII, Japan)
- Jung-Sing Jwo (Tunghai University, Taiwan)
- Rick Kazman (University of Hawaii, USA)
- Foutse Khomh (Polytechnique Montreal, Canada)
- Takashi Kobayashi (Institute of Science Tokyo, Japan)
- Judith Michael (University of Regensburg, Germany)
- Henry Muccini (University of L'Aquila, Italy)
- Cesare Pautasso (University of Lugano, Switzerland)
- Abhik Roychoudhury (National University of Singapore, Singapore)
- Marija Slavkovik (University of Bergen, Norway)
- Hironori Takeuchi (Musashi University, Japan)
- Catia Trubiani (Gran Sasso Science Institute, Italy)
- Naoyasu Ubayashi (Waseda University, Japan)
- Hironori Washizaki (Waseda University, Japan)
- Sherry (Xiwei) Xu (CSIRO, Australia)
- Joseph Yoder (The Refactory, USA)

Group photos

Yosegaki

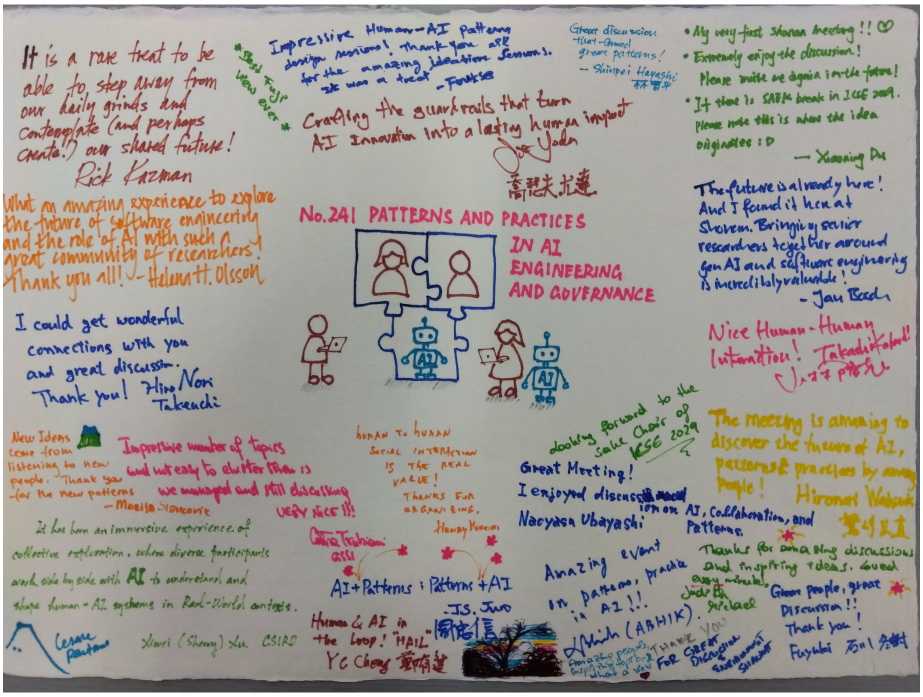

Figure 2: Yosegaki