



# Rethinking Software Engineering in the Era of Foundation Models

A Curated Catalogue of Challenges in the Development of Trustworthy FMware

Ahmed E. Hassan  
Queen's University  
Kingston, Canada

Dayi Lin  
Gopi Krishnan Rajbahadur  
Keheliya Gallaba  
Huawei Canada

Filipe Roseiro Cogo  
Boyuan Chen  
Haoxiang Zhang  
Huawei Canada

Kishanthan Thangarajah  
Gustavo Oliva  
Huawei Canada

Jiahuei (Justina) Lin  
Wali Mohammad Abdullah  
Huawei Canada

Zhen Ming (Jack) Jiang  
York University  
Canada

## ABSTRACT

Foundation models (FMs), such as Large Language Models (LLMs), have revolutionized software development by enabling new use cases and business models. We refer to software built using FMs as FMware. The unique properties of FMware (e.g., prompts, agents and the need for orchestration), coupled with the intrinsic limitations of FMs (e.g., hallucination) lead to a completely new set of software engineering challenges. Based on our industrial experience, we identified ten key SE4FMware challenges that have caused enterprise FMware development to be unproductive, costly, and risky. For each of those challenges, we state the path for innovation that we envision. We hope that the disclosure of the challenges will not only raise awareness but also promote deeper and further discussions, knowledge sharing, and innovative solutions.

## CCS CONCEPTS

• **Software and its engineering** → **Software creation and management**.

## KEYWORDS

Foundation models, Large Language Models, AIware, FMware

### ACM Reference Format:

Ahmed E. Hassan, Dayi Lin, Gopi Krishnan Rajbahadur, Keheliya Gallaba, Filipe Roseiro Cogo, Boyuan Chen, Haoxiang Zhang, Kishanthan Thangarajah, Gustavo Oliva, Jiahuei (Justina) Lin, Wali Mohammad Abdullah, and Zhen Ming (Jack) Jiang. 2024. Rethinking Software Engineering in the Era of Foundation Models: A Curated Catalogue of Challenges in the Development of Trustworthy FMware. In *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering (FSE Companion '24)*, July 15–19, 2024, Porto de Galinhas, Brazil. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3663529.3663849>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*FSE Companion '24, July 15–19, 2024, Porto de Galinhas, Brazil*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0658-5/24/07

<https://doi.org/10.1145/3663529.3663849>

## 1 INTRODUCTION

FMware refers to the type of software that uses foundation models (FMs), such as Large Language Models (LLMs), as one of its building blocks. Due to their general capabilities, flexibility, and natural language interface, FMs open doors to a new paradigm of software development where domain experts without deep programming or AI skills are empowered to build FMware for a variety of use cases [13]. The market size of FMware is estimated to grow at a compound annual growth rate of 35.9% from 2024 to 2030 [37].

Despite the benefits of FMs, building trustworthy FMware remains considerably difficult. We observe that developers constantly suffer from three key pain points when integrating FMs in software systems: (i) **low productivity** throughout the lifecycle, due to the lack of mature practices and integrated tooling, (ii) **high risk**, due to the stochastic and nondeterministic nature of FMs, and (iii) **high operational cost**, due to the high resource consumption of FMs. And we are not alone. In a recent article by Microsoft [79], 26 professional software engineers responsible for building Copilot-like products echo our pain points: prompt engineering and testing are extremely time-consuming and resource-constrained. Additionally, the participants were frustrated with the lack of comprehensive tooling and best practices in this nascent domain.

In this paper, we discuss a curated catalogue of software engineering (SE) challenges for developing FMware (i.e., SE4FMware). These challenges were identified based on (i) surveys of academic and grey literature, (ii) in-depth discussions with industrial and academic leaders (e.g., during SEMLA 2023 [51] and the FM+SE Vision 2030 [41] event with over 100 attendees from many leading companies), (iii) meetings with our customers and our development teams to understand the functional and non-functional needs of their FMware, and (iv) our practical experience in designing, implementing, and maintaining our in-house FMware lifecycle engineering platform (FMArts) and several complex FMware for strategic customers. For each challenge, we discuss the innovation path that we envision. These innovation paths call for significant breakthroughs in software engineering research and practice.

## 2 SOFTWARE ENGINEERING FOR FMWARE

In this section, we first describe the software generation of FMware and how it relates to prior ones (Section 2.1). Then, we describe the inherent limitations of FMware (Section 2.2). Finally, we present our view of the FMware development lifecycle (Section 2.3).

## 2.1 FMware in Context

Over the past decade, we have witnessed the birth and evolution of several generations of software. Each new generation is characterized by drastic changes to the development paradigm, often involving new lifecycles, core assets, and roles. *Codeware* (a.k.a. Software 1.0) development involves professional programmers writing logic-based source code, following a typical development lifecycle that includes a well-established set of phases (requirements, design, implementation, testing, deployment, and maintenance) that are carried out either sequentially (e.g., the waterfall method) or iteratively (e.g., the agile family of development processes). As machine learning models started to be embedded in software in recent years, AI experts started to drive the development of a new generation of software known as *Neuralware* (a.k.a. Software 2.0) by constructing a new set of assets such as *datasets* and *models* in a significantly more iterative and experiment-driven lifecycle [6]. Since most of the traditional SE practices and tools were no longer applicable for Neuralware, many research efforts have been made in the SE community to address the unique challenges of Neuralware throughout its lifecycle (e.g., versioning of AI models [49]).

The rapid development and adoption of foundation models has led to the emergence of yet another generation of software known as *Promptware* (a.k.a. Software 3.0). Unlike prior generations, Promptware is developed through *prompts* written in natural language. Such a paradigm shift unlocks the potential for software to be developed by *software makers*, i.e., people without deep programming and AI skills, ultimately democratizing software creation. ChatGPT is a popular example of Promptware.

Recently, people have started to leverage the reasoning capabilities of FMs to implement agents that are capable of decision-making, taking actions in an environment, and even interacting with other agents. In particular, agents achieve their goals by leveraging all prior software generations. We call this newest generation *Agentware* (a.k.a. Software 4.0). AutoGPT [1], BabyAGI [2], CrewAI [3], and MiniAGI [4] are few examples of this recent development.

Finally, we foresee the birth of a new software generation once artificial general intelligence (AGI) becomes a reality. We refer to this generation as *Mindware* (Software 5.0). Mindware consists of a paradigm where development is led mostly by autonomous agents. Humans themselves are expected to intervene less in software development. Instead, humans will focus on creating and refining *constitutions* [10] while overseeing the agents [57].

We collectively refer to Promptware and Agentware as *FMware* (due to the key role of FMs in those software generations). Similarly, we refer to Neuralware, Promptware, Agentware, and Mindware as *AIware*. We also note that latter generations of software could include components from prior generations, and will coexist with earlier generations rather than replace them.

## 2.2 FMware and the Intrinsic Limitations of FMs

Despite the huge market potential of FMware, FMs exhibit several *intrinsic limitations* that complicate the development of FMware:

- **Complex Task Limitations:** FMs are unable to handle complex multi-step tasks that require planning [93]. FMs also have limited reasoning capabilities [103] unless prompted in particular ways (e.g., using prompt engineering techniques such as chain-of-thought prompting [96]). Finally, FMs exhibit uncontrollable behavior in

multi-round interactions due to the limited size of the context window and may get confused with long prompts despite growing context window size [61].

- **Hallucination Limitations:** FMs are prone to hallucination and can confidently generate factually wrong responses [46].

- **Closed Loop Limitations:** FMs by themselves are unable to interact with external resources [107].

## 2.3 The FMware Development Lifecycle

The distinguishing characteristics of FMware (Section 2.1) and the intrinsic limitations of FMs (Section 2.2) lead to an entirely novel software development lifecycle. We summarize our view on the lifecycle of FMware in Figure 1. Such a view was consolidated based on our literature survey, meetings with our customers and development teams, and our practical experience in building both complex FMware as well as our own FMware lifecycle engineering platform. Yet, we acknowledge that this view may not be complete due to the pace at which the field evolves. Nevertheless, we hope that our concrete view will kick start urgently needed discussions around the FMware lifecycle, including the naming and scope of the various phases and activities.

**Requirements engineering.** The process starts at the *requirements engineering* phase. While requirements engineering should preserve some degree of technology agnosticism (i.e., focus more on the *what* than the *how*), we observe that a lack of understanding of the capabilities and drawbacks of fundamental technologies in the solution space (e.g., FMs) often leads to problems (e.g., unrealistic specification of system behavior). Therefore, the elicitation and specification of requirements should account for these important characteristics (e.g., FMs lack deterministic behavior).

**Cognitive architecture planning.** After requirements are collected, FMware goes into a new phase that we call *cognitive architecture planning*. This phase entails choosing the coordination mechanism to dictate what software systems (possibly from different generations) will be invoked, how, and in which order. In our mental model of cognitive architectures, the coordination mechanism choice draws the line between Promptware and Agentware. In Promptware, the coordination mechanism is a *static workflow* that prescribes a set of tasks and the order in which those will be executed (and a task can always be another workflow). These workflows are also known simply as *flows* (e.g., in Microsoft Promptflow [70]). We refer to the task of producing workflows as *workflow engineering* (see Promptware in Figure 1). While humans might leverage FMs to help them design (or co-design) these workflows, the goal is still to obtain a static workflow. In the case of Agentware, the coordination mechanism is driven by autonomous agents. That is, agents decide the steps **and** the order in which those will be taken. In *multi-agent* cognitive architectures, coordination is achieved as a result of inter-agent interactions (e.g., Microsoft Autogen [100]). We refer to the task of designing agents (and their communication protocol in the case of multi-agent architectures) as *agentification* (see Agentware in Figure 1). In both Promptware and Agentware, task execution can always be accomplished by (or with the help of an) agent. It is also worth noting that although agents have higher autonomy, a certain level of human intervention is typically needed to reduce risks in FMware. For instance,

**Table 1: Cognitive architecture patterns.**

Type of FMware	Architectural Pattern	Decide Output of Step	Decide Which Steps to Take	Determine What Sequence of Steps Are Available
Promptware	FM Call	FM (one step only)	Workflow	Workflow
Promptware	FM Chain	FM (multiple steps)	Workflow	Workflow
Agentware	Router Agent	Agent	Agent	Workflow
Agentware	State Machine Agent	Agent	Agent	Workflow
Agentware	Autonomous Agent	Agent	Agent	Agent

*constitutions* [10] might be used to ensure that agents adhere to certain basic principles and ethical standards.

Architectural patterns [32] are starting to emerge for *cognitive architectures*. LangChain recently published a list of such patterns [56], which we adapt to our terminology and show in Table 1. The patterns are listed in order of controllability. For instance, in the first (topmost) pattern, the steps to be taken and their order are already established via a static workflow, thus giving humans more control. In the last pattern, the steps to be taken and their order are dynamically determined by autonomous agents, giving humans little control (but more room for emergent behaviors and reasoning).

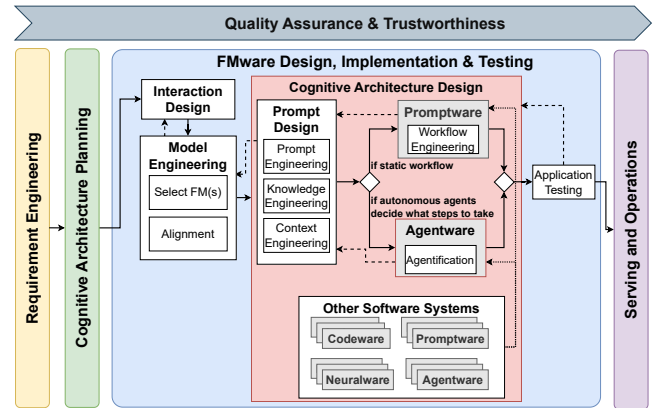
As the field matures and different architectures are deployed in enterprise settings, we believe new cognitive architecture patterns will emerge, and their trade-offs will become clearer.

**FMware design, implementation & testing.** After the cognitive architecture is planned, the *design, implementation & testing* phase starts. This phase tends to be highly iterative (note the dotted back arrows). In our experience, developers typically jump between steps as the design evolves and requirements become clearer.

The *design, implementation & testing* phase starts with *interaction design* [81], which involves designing how end-users will interface with the system. In the context of FMware, many prescribe a back-and-forth conversation between the end-users and the system (e.g., ChatGPT). For Agentware, when and how would users provide oversight to agent behavior should be considered here.

Next, we advance to the (*foundation*) *model engineering* task, which comprises selecting one or more FMs and aligning them if necessary. *Model alignment* is the process of modifying the parameters of an FM using a set of techniques and carefully crafted data, such that the resulting FMware is aligned with the requirements. Examples of alignment techniques include supervised fine-tuning [16] and reinforcement learning from human feedback (RLHF) [76].

Subsequently, we jump to the *cognitive architecture design* sub-phase. This subphase starts with the *prompt design* step. In this step, prompts are designed according to the previously chosen cognitive architecture. Prompt design encompasses *prompt engineering*, *knowledge engineering*, and *context engineering*. Prompts are written using prompt engineering techniques that maximize the chances that the FM will consistently provide the desired response. Context engineering entails techniques for manipulating and enriching the context of prompts (e.g., prior interaction history and quick patches of FMware behavior). Knowledge engineering entails creating, managing, and retrieving domain knowledge to ground the FMware. For example, grounding techniques such as retrieval-augmented generation (RAG) [58] have proven useful in mitigating FM hallucination. Once prompts are designed, the lifecycle moves to either the *workflow engineering* (in the case where coordination is dictated by static workflows) or the *agentification* step (in the case where coordination is dictated by agents). In our experience, prompt

**Figure 1: Our view of the FMware development lifecycle.**

design and workflow engineering/agentification form a highly iterative loop, as one step drives and refines the other. As discussed above, developers might leverage cognitive architectural patterns to implement the chosen type of coordination. In complex FMware, the workflow engineering/agentification step typically involves invoking several software systems from different generations. For instance, an FMware could use a Python interpreter (Codeware) to execute Python code, a neural network (Neuralware) to recognize patterns in images, an interactive copyeditor to support documentation (Promptware), and a general task planner (Agentware).

Finally, the lifecycle advances to *application testing*. In this step, the whole FMware must be tested. Testing is particularly hard due to the non-determinism of FMs as well as the presence of heterogeneous multi-generational software (i.e., FMware is ultimately a complex *system of systems*).

**Serving and operations.** The last phase is *servicing and operations*. This phase concerns the packaging, release, deployment, telemetry, monitoring, observation, and patching of FMware. Since the quality of FMware output is hard to quantify (see Section 3.9), traditional observability must be expanded to include the collection, processing, and analysis of data that represent end-user intent and experience, including explicit user rating of the FMware output, or implicit behavior such as correction of the output. We refer to this new type of observability as *semantic observability*.

**Quality assurance (QA) & trustworthiness.** The intrinsic characteristics of FMs (e.g., non-determinism), as well as their limitations (Section 2.2), pose several challenges to QA and trustworthiness, which are cross-cutting and thus are enforced throughout the entire lifecycle (note the upper arrow in Figure 1).

Overall, the unique characteristics of the FMware development lifecycle demand the research community to study and design SE principles and best practices for developing trustworthy FMware. We refer to this new discipline as SE4FMware.

### 3 CHALLENGES IN ENGINEERING PRODUCTIONIZED FMWARE

In this section, we introduce ten SE4FMware challenges. As mentioned in Section 1, these challenges were identified based on (i) literature reviews, (ii) in-depth discussions with industrial and academic leaders, (iii) meetings with our customers to understand the

requirements of their applications, and (iv) our practical experience in designing, implementing, and maintaining complex FMware systems for those customers.

To accomplish item (iv), our team created the FMArts platform as our long-term endeavor to design and implement a full-stack lifecycle engineering platform for the development of trustworthy FMware. As such, the challenges that we discuss in this section are heavily influenced by our experience in developing our platform.

Each challenge is tagged in blue with the FMware development lifecycle phase to which it belongs. Additionally, each challenge is tagged in red according to the pain point(s) that it causes, namely *low productivity*, *high risk*, and *high operational cost*.

### 3.1 Managing Alignment Data

🔧 Design, Dev. & Test   ● Productivity   ● Operational Cost

**Description.** To produce well-aligned models for trustworthy FMware, an alignment approach that seamlessly combines automation with strategic human input is essential. Pre-trained FMs often do not produce outputs aligned with the requirements and goals of enterprise use cases out of the box. Therefore, aligning an FM with specialized enterprise data is essential to meet requirements and achieve market differentiation. Similar to Neuralware shifting development paradigm from code-driven to data-driven, aligning a model is achieved through data-oriented development and needs to cope with challenges that are associated with data labelling, data debugging, data testing, and data versioning for FMware. Although Neuralware also requires human effort to manage training datasets, aligning FMs involves new types of data such as instructions, reference responses in natural language, ratings, and rankings of potential model outputs, requiring different efforts in the curation and maintenance processes. Another unique challenge for FMware is determining if a particular piece of data is part of the FM's training data and avoiding problems such as data leakage during evaluation. Consistent with recent industry reports [17], we observe that such processes demand significant time and financial resources. In addition, it is challenging to accurately and quantitatively gauge the adequate amount of high-quality alignment data for a specific use case, creating challenges for budgeting and time management.

**State of practice.** Researchers and practitioners are turning to active learning [92, 110] and weak supervision methods [9, 12] to reduce the involved effort in constructing alignment data. For example, Snorkel AI [85] employs data programming to generate labels using expert knowledge, while the Self-Instruct framework [95], a pioneering work in the area of auto-labeling, uses a powerful FM (e.g., GPT-4) to automatically generate labeled datasets from seed prompts, bypassing the manual labeling process.

Collecting alignment data in FMware is still too expensive and slow. Active learning methods demand significant manual input, with over 50% of the data needing manual labels [67], which is often unfeasible. Additionally, removing human labels can harm model alignment, as relying solely on automated labeling can lead to performance degradation, catastrophic forgetting, and model collapse [63, 90]. Another line of research attempts to reduce the needed amount of alignment data by focusing on the quality of the alignment data. Two pioneering works in this area [39, 113] show that fine-tuning an FM on textbook-quality alignment data results in

an FM needing only a small amount of data for alignment. However, these methods still require a lot of manual effort to identify high-quality data. A recent effort by Zhang et al. [111] proposes an automatic method to identify a subset of high-quality alignment data within a larger dataset [39].

To support data debugging, researchers have employed statistical techniques to identify, select, and exclude undesired (e.g., biased) features, labels, and data points from training sets while ensuring the models' performance is on par (or better) than its counterpart trained in the original data [18, 59]. In the realm of FMs, influence functions and feature importance analysis can identify training data points that significantly impact the model's predictions [38, 88]. In addition, tools for data testing [14] either support the iterative process of cleaning up data while assessing its impact on models' performance [53] or automate such process [52]. However, such tools still need to account for the particularities of FMware, such as the lack of structure and the high dimensionality of the alignment data, including the opaqueness of FMs, which makes it difficult to determine the direct impact of training data on the model's outputs.

**Innovation path.** The management of alignment data for FMware calls for data IDEs that can fully support the model alignment lifecycle, including labeling, debugging, testing, and versioning data. Specific directions for such data IDEs include the integration of model- and human-in-the-loop, the programmatic labeling (data programming) and identification of inconsistencies or biases in the data, and features to support collaboration, such as supporting the aggregation and review of labeled data by different labelers. In our FMArts platform, we address some of these requirements by offering an *alignment data studio* that enables the curation of alignment data with an automated labeling approach that requires very little manual intervention. We also provide robust review mechanisms with support for debugging the labeled data. However, identifying biases, ensuring compliance, and testing the created alignment data remain a work in progress.

In addition, open and inner source data offer a significant window of opportunities to model alignment. Still, data may have licensing restrictions limiting their use, particularly in commercial applications. Therefore, such data IDEs should provide asset management capabilities for open and inner source data to reduce risks, for instance, by ensuring that the data is accurate, consistent, complete, and representative of a task. Also, as software makers and their domain expertise (Section 2.1) play a major role in data programming for FMware (e.g., by defining accurate labeling functions), data IDEs need to embrace the characteristics of its users, such as their potential lack of formal SE training.

### 3.2 Crafting Effective Prompts

🔧 Design, Dev. & Test   ● Productivity

**Description.** FMware needs an approach to prompt engineering that reduces human involvement and improves development efficiency and application quality, as well as prompt transferability and fragility. As the communication interface for FM, prompts at their current form are too low-level (i.e., at the level of input tokens to the FM) and fragile. This makes engineering prompts a non-trivial task due to the time-consuming and non-intuitive nature of iteratively hand-writing prompts [91]. This trial-and-error

approach *lacks a clear structure* and *imposes a productivity bottleneck* on prompt engineering workflows, as multiple iterations of a prompt are created before the most suitable one for a given task is identified, particularly for complex functionalities that require chaining multiple prompts [101, 102]. To exacerbate the situation, *prompts are fragile*, in the sense that slight variations in a prompt can cause a significant change in the generated outputs [33], leading to continuous prompt adaptation and reevaluation, which might be costly if not properly managed. Similarly, small changes to the underlying FM or its hyper-parameters render existing prompts ineffective, ultimately hindering the generality of an implemented FMware solution. For example, as OpenAI constantly updates its served models [5], preliminary evidence shows that, when the same FM (e.g., GPT-4) is inferred with the same prompt, the generated results by the FM differ at two different time instances [21].

Due to the unpredictable behavior of FMs, ensuring the accuracy and reliability of outputs becomes a significant challenge, as FMware developers must navigate the complex landscape of model behaviors without clear indicators of how undesired responses are being generated. Such lack of transparency and control makes achieving consistent and dependable results from FMware difficult.

**State of practice.** Visual tools [91] (akin to playgrounds) attempt to assist FMware developers by providing inference transparency (e.g., Prompt IDE [104] by xAI allows FMware developers to observe the generation process of FMs, such as sampling probabilities) and token-level explainability (e.g., Aleph Alpha [25] offers token-level attribution between input tokens and output tokens). While these capabilities are certainly helpful, they do not support higher level prompt debugging, i.e., identifying the logical units in a prompt lead to suboptimal results (e.g., a misleading few-shot example).

Moreover, researchers and practitioners are actively looking for approaches to provide a cohesive and expressive framework for automating and supporting prompt engineering without forcing FMware developers to rely on completely manual, unstructured, or error-prone processes [62]. Several research endeavours attempt to fully [43] or semi-automate [44, 106] prompt engineering [109, 115] through *search-based approaches* or *reinforcement learning-based approaches* [26]. For example, Automatic Prompt Engineer [115] generates several instruction-based prompt candidates via a recursive process, and PromptBreeder [29] uses a genetic algorithm to select the most appropriate prompt. DSPy [50] combines a programming model with an optimization compiler to generate the best combination of prompts for a model. However, most approaches optimize prompts at the token level in isolation from other factors that might impact the results, such as model candidates, inference parameters, and hardware environments. Yet, in our industrial setting, we see FMware developers experimenting with a large set of FMs and fine-tuned variants, as well as prompt candidates under different settings to find the most suitable combinations.

**Innovation path.** To improve the quality of FMware output, researchers and practitioners have optimized the interaction with FMs in two directions: (i) breaking down one prompt for a complex task to prompting patterns that involve a series of prompts (i.e., prompt chains), sometimes with complex branching and conditions (e.g., chain-of-thought and tree-of-thought) and (ii) human

involvement inside a chain of prompts to provide feedback to intermediate outputs (e.g., for code generation, one could first prompt the FM to generate a description of the algorithm, then involve humans to correct the description, and prompt the FM again to implement the algorithm as code). Instead of focusing exclusively on prompt engineering techniques, researchers should also explore techniques for the design, debugging, visualization, and evolution of complex cognitive architectures that prescribe different levels of human involvement. Such techniques should incorporate higher-level prompts to support the dynamic selection and adjustment of lower-level prompts to sustain meaningful interactions based on the model's responses, context, and end-user interactions.

Designing complex cognitive architectures requires a proper *prompt IDE*. This IDE should support real-time collaborative features. Also, mechanisms to deal with just-in-time quality assurance should be provided, such as prompt debuggers and linters. In certain cognitive architectures, it is also important for prompt results to be reproducible. Yet, simply setting the FM temperature to zero is not enough in certain use cases. Hence, research must be carried out to create mechanisms that enforce full reproducibility [105]. Finally, to foster a rich and open ecosystem of functionalities, prompt IDEs should offer a plug-and-play architecture that facilitates integration with community-contributed tools and services (similar to leading SE ecosystems like GitHub and Eclipse).

To fulfill these requirements, the *prompt studio* in our FMArts platform provides an environment where version-controlled prompts can be connected to different FMs. This connectivity facilitates prompt refinement in conjunction with various FMs and execution setting configurations, including near real-time visualization of prompt optimization results. Our prompt studio offers asset management capabilities that enable prompts to be collaboratively developed, reviewed, and shared either globally (at the platform's scope) or locally (at a project's scope, using role-based access control), thus facilitating teamwork. Our prompt studio also incorporates debugging capabilities, such as prompt interpretation tools that dissect the contributions of individual prompt components to the generated output. Token-level explanation is also supported for finer-grain debugging. FMArts also provides a *knowledge studio*, to handle the upload, tokenization, vectorization, and retrieval of grounding data to be used in prompts, ultimately contributing to quality assurance (e.g., mitigation of model hallucination).

### 3.3 Multi-generational Software

⚙️ Design, Dev. & Test   ● Productivity

**Description.** FMware developers face the intricate task of integrating different generations of legacy software systems, with each having its own data models, communication protocols and execution processes. This integration process demands supplemental effort to ensure seamless interoperability of many legacy software systems, thus hindering developer productivity. To maximize the usage of legacy software systems across multiple FMware and avoid redundant work, FMware developers need to either standardize the data formats of these legacy software systems or transform the data models at the post-retrieval step. However, such an integration approach introduces a significant workload. In addition, changes in the shared legacy software systems impact the downstream FMware,

as FMs are sensitive to prompt changes (see Section 3.2). As such, FMware developers may opt to customize the legacy software for an individual FMware, facing maintenance overhead.

**State of practice.** To integrate multi-generational software, the current practices either use the concept of plugins to construct reusable units or package legacy systems as tools that FMware could reach. Semantic Kernel [69] leverages plugins through the construct of *native functions* to encapsulate the capabilities of a legacy system into functionality that can then be run by the kernel. Approaches such as LlamaIndex [65], ReAct [108], TaskMatrix [60], Toolformer [87], and ToolLLM [83] assist FMware developers to interact with legacy systems. However, the use of such techniques is often limited by the additional effort that is required from FMware developers to repackaging existing tools or finetune FMs to be compatible.

**Innovation path.** Many current research efforts and frameworks assume FMware as a green-field project. Integrating with the large amounts of legacy software already in existence is not well addressed. The software engineering community should take into account how to effectively integrate FMware with legacy systems by considering *reusability* and *modularity* principles. In our FMarts platform, *cross-generational agents* are designed to reduce the effort of integrating multiple software generations in FMware, as well as to address evolution and compatibility concerns. This design objective is achieved by encapsulating behavior that is common to each of the software generations into different modules (which we call *skills*) and then unifying the interface of such modules into agents. Agents are designed using our *agent studio*.

### 3.4 Degree of Controllability

⚙️ Design, Dev. & Test   ⚙️ Serving & Ops   🚫 Risk

**Description.** Trustworthy software systems, especially in sensitive domains such as finance, healthcare, and manufacturing, require robust controllability to prevent the possibility of misbehaviors causing unintended catastrophic consequences [97]. of FMs. The potential of systems that exert direct control over the world to cause harm in a way that humans cannot necessarily correct or oversee has been identified as a concrete problem in AI safety [7].

However, given the inherent stochastic nature of FMs (Section 2), determining the degree of controllability of FMware is challenging. It is critical to be mindful of the control of an FMware throughout its lifecycle. Specifically, we discuss controllability in three dimensions: (1) how, (2) when, and (3) who. FMware developers plan the control mechanisms of an FMware during the *cognitive architecture planning* and *cognitive architecture design* phases, both of which happen after requirements are elicited and analyzed. The *how* dimension includes determining the data granularity to protect sensitive data, access to various tools or services within FMware, and the autonomy to execute behaviors with or without human intervention. It is challenging to define the *how* aspect as FMware achieve certain tasks with different complexity. In terms of the *when* dimension, controllability of FMware varies across its lifecycle. FMware developers calibrate the degree of control of a given FMware over time and may disable certain control features if the FMware behaves unexpectedly in production. Lastly, the *who* dimension describes which stakeholders are authorized to grant control permissions, including FMware developers, end-users or other roles (e.g., admin).

For example, during the execution of FMware, end-users grant permission to authorize an FMware to execute sub-tasks to achieve the high-level goals specified in the original plan.

**State of practice.** Restricted approaches, such as standard operating procedures (SOPs) and standard work instructions (SWIs), are employed for humans in industries where safety and quality are critical. Given the autonomous capability of FMware, there is an opportunity to make use of such existing approaches to effectively control FMware. Several frameworks have been proposed to customize the degree of control. Copilot Workspace [35] provides control by communicating the tasks that will be done in advance to the user so they can decide to continue on that path of action. AutoGen [100] is another framework that enables FMware development and offers mechanisms to solicit human inputs to offer different levels of control. Zhou et al. [114] propose to build controllable agents using symbolic planning. MetaGPT [42] leverages the assembly line paradigm and encodes SOPs into prompts to enhance human-encoded and FM-powered coordination, and thus controllability. All of these existing efforts and our own experience, highlight the importance of maintaining effective human oversight as a key requirement when supporting FMware.

**Innovation Path.** Fine-grained boundaries are defined in classic software systems that are designed to be extensible by third-party components (e.g., application permission model in Android [36] and the extensions architecture of Google Chrome [36]) when accessing resources to preserve the privacy and security of end-users. If elevated privileges are needed, components can explicitly request for permission from end-users. Similarly, FMware should leverage such a concept of boundaries to develop mechanisms to enforce constraints at a fine-grained level.

Existing concepts in software engineering, including *modularity* and *single responsibility principle*, are applicable to decompose FMware functionality into small and independent components (e.g., agents) with well-defined interfaces. Therefore, component behaviours can be controlled in a granular manner. The need for abstractions to define control constraints facilitates the creation of consumable agents and tools for FMware developers. Our FMarts platform offers an *orchestration studio* for the creation of *SOP-based workflows* that constrain multi-agent behavior and interactions. Moreover, our platform auto-generates an architecture view by analyzing all assets in one FMware application. Developers can use this view to recognize deviations in implementation from the intended design. This holistic view aids in managing design principles such as *cohesion*, *coupling*, and *separation of concerns* within an FMware.

### 3.5 Compliance & Regulations

⚙️ Requirements Eng.   ⚙️ Design, Dev. & Test   ⚙️ Serving & Ops   🚫 Risk

**Description.** Trustworthy software systems are subject to stringent regulatory compliance requirements to ensure their security, safety, and compliance with relevant laws and regulations, while FMware is increasingly being used in decision-making processes where it is necessary to comply with such requirements. Despite governments worldwide having recently released regulations for FMs and FMware, it is still challenging to fully comply with them [40]. An

example is the EU AI Act [28], which, in its Article 14, calls for high-risk FMware to be transparent and capable of human oversight, despite this requirement being currently intractable for FMware.

A particularly relevant regulation is the GDPR in the European Union, which requires that personal and high-risk data should always be protected and should not leave secure bounds. For FMware that uses FMs as a third-party inference service, the possibility of data leaking beyond the secure enterprise bounds is high [79]. Several large enterprises like Microsoft and GitHub resort to self-hosting FMs to avoid such issues. However, not all enterprises have the means to self-host such large models. Therefore, SE4FMware needs to put forward measures, techniques, and tools to ensure regulatory compliance in such settings.

Ensuring licensing compliance within a FMware system is notably challenged by the opaque nature of FM development. FM providers do not always disclose their data sources [66, 89] or the usage of protected [19] and personal data [71], making it difficult for some FMware to comply with various data, code licenses, and copyright restrictions. Moreover, the necessity for diverse datasets and the concurrent operation of multiple FMs introduce significant complexities in license management, particularly when faced with ambiguous, conflicting terms [11, 84]. Notably, prominent FMware implementations like ChatGPT and GitHub Copilot have encountered legal challenges related to these compliance issues [23].

The high costs and extensive resources required to train FMs push enterprises towards utilizing open-source FMs. A prime example is LLaMA, an FM that was acclaimed for its open-source status when first released, yet restricting commercial usage and revoking licenses when exceeding 700 million subscribers. This issue is further compounded in the FMware development sphere by the increasing prevalence of non-Open Source Initiative (OSI) approved licenses for both FMs and datasets, undermining traditional open-source values and intensifying compliance dilemmas.

**State of practice.** Efforts towards ensuring regulatory compliance of FMware remain largely manual and process-heavy. Brajovic et al. [15] propose to enhance the model and data cards that are currently used to document FMs with more relevant fields that could enable compliance with the EU AI Act. In addition, they propose two new schemas to document the relevant details of the FM's use case and deployment settings. However, they do not address how such reporting can be done for a complex FMware. As for data compliance, FMware can use a method that identifies copyright or license-protected data in FMs by analyzing the fingerprint of output tokens with those of protected data and FM outputs [89]. Another approach taken by several government bodies [74] and research initiatives [75] is to provide risk management frameworks that are applicable to FMware, which could, in turn, help ensure compliance with specific regulations.

License compliance efforts often concentrate on the individual FM or dataset, overlooking the intricacies of compliance within the FMware domain. A notable development in this area is the introduction of Responsible AI Licenses (RAIL) [24], aimed at incorporating behavioral and use-based restrictions into AI technology agreements to address ethical and regulatory challenges. Specifically, the RAIL-M licenses seek to embed AI-specific ethical standards through copyright licensing. Additionally, the Open

Data Commons Public Domain Dedication and License (PDDL) and Creative Commons (CC) licenses strive to offer similar solutions for datasets [20]. Different from the licensing efforts, open source projects like OpenDataology [99] seek to decompose the dataset license into machine-readable AI-friendly use-case scenarios to enable better license interaction analysis. Despite these advancements, certain critical issues remain unaddressed, impeding comprehensive FMware compliance. These include focusing on singular models or datasets rather than integrated FMware systems, disregarding license interaction problems [20, 84], and overlooking the scenarios involving lifelong-evolving agents (e.g., learning a new behavior could violate that agent's current license).

Towards enabling FMware level compliance, recent initiatives to enhance Software Bill of Materials (BOMs) to represent AI and datasets (e.g., SPDX [31]) offer a promising method for managing FMware and its dependencies (e.g., alignment data) in a format conducive to automated license compliance and vulnerability assessment. The advent of AIBOMs could facilitate the automatic risk analysis for FMware.

**Innovation path.** Automatically determining FMware's composition and producing machine-readable SBOMs to catalogue FMs (FMBOMs) as well as every asset of an FMware (FMwareBOMs) is essential. Techniques such as formal verification can also be used to analyze these FMwareBOMs and the FMware's output to ensure an FMware's regulatory and license compliance. In FMarts, we implemented the generation of a preliminary version of FMwareBOM closely modeled after the SPDX 3.0's AI and Dataset profile [31]. With the help of these FMwareBOMs, FMarts conducts preliminary compliance analysis with enterprise rules and regulations using *automated reasoning* [86]. However, formal verification for FMwareBOMs remains an active area of development.

To mitigate the risk of legal consequences, FMware-specific licenses that consider the interplay between FM, dataset, code, and agent licenses are essential. Implementing tools akin to Fossilogy [98] tailored for FMware is pivotal in supporting enterprise FMware to remain compliant. Furthermore, documentation standards for disclosing compliance verification procedures and the shortcomings for FMware need to be established.

Another fundamental endeavor is the design of robust mechanisms to ensure enterprise and personal data privacy in FMware. Safeguard mechanisms should be systematically implemented to ensure that the incoming and outgoing data comply with regulations at all levels. In addition, despite telemetry information (e.g., logs) being of fundamental importance to understand FMware's behavior (including if the system is fully compliant), balancing the amount of information that is revealed to FMware operators is still an ongoing challenge. Privacy-preserving techniques can be applied so that operators get the most out of the telemetry data while preventing sensitive information disclosure, such as entire prompt details [79].

### 3.6 Limited Collaboration Support

🔧 Design, Dev. & Test    🚀 Productivity

**Description.** For complex software, development is almost always collaborative. However, there is currently very limited support for the collaborative development of FMware. Such a limitation severely hampers team productivity, inner-source, open-source,

and FMware ecosystem growth. Tackling this limitation involves rethinking the enablers of collaboration: from version control all the way to hubs (sharing platforms). In particular, the following points need to be addressed:

- *What to version and how.* In Codeware, we know that source code must be version controlled and the enabling technologies (e.g., Git) are very mature at this point. However, FMware introduces new types of assets. For instance, prompts are typically described using natural language and stored as raw text. As a consequence, logical units of a prompt (e.g., examples) across different prompts become a new type of reusable unit. In an effort to promote reuse, prompts are also sometimes written and stored as templates (e.g., using the Jinja2 template engine [82]). However, due to the prompt fragility problem (see Section 3.2), very similar prompts could yield very different results. Such a problem calls for version controlling all iterations where placeholders have been replaced. However, in that case, the connection between the template and its variants (as well as variants of variants) becomes lost, since hierarchical information is not natively preserved. More generally, the granularity at which prompts are represented, stored, and version-controlled must be carefully explored. Version controlling agents is also difficult due to the inherent characteristics of agents [94]. First, agents are often implemented as an amalgamation of source code, FMs, and potentially other assets. This brings up the question of how exactly agents should be serialized and version controlled. Moreover, intelligent agents are expected to incrementally acquire, update, accumulate, and exploit knowledge throughout their lifetime (a.k.a *continual learning* and *incremental learning*) [94]. In such a context where agents learn over time, as well as forget things that they had previously learned, it remains unclear *when* and *how* to version control the agent (e.g., when should snapshots be captured?).

- *What formats and standards to use.* The lack of standards and protocols for FMware assets significantly hamper interoperability, i.e., the ability to reuse assets created by different organizations. Additionally, the lack of standards and protocols limits the ability to design development tools that work with a wide range of assets (e.g., agents) out-of-the-box. As discussed by the AI Engineer Foundation [30], developers are building agents in their own way, making communication with those agents very hard since their interface is different every time.

- *How to package and distribute.* While significant advancements have been made in the packaging and distribution of AI models (e.g., via Hugging Face), it remains unclear what the stereotypical release of an agent or a prompt should look like. For instance, models in Hugging Face are versioned and include the datasets on which they were trained. Given the adaptive nature of agents, how should the “training data” (e.g., environment definition, actions on the environment, signals received from the environment, interactions with other agents and humans) be described and released? More generally, how does one package an agent? These release engineering problems need to be tackled to enable widespread sharing of assets and the future establishment of an ecosystem of FMware assets and creators.

**State of practice.** Prompts and prompt templates are typically recorded in Git. Simpler agents that encode a list of sequential steps are typically described using a markup language such as YAML

(e.g., flows in Microsoft PromptFlow [70]). There are no established standards for describing prompts and agents. In the space of agents, one of the most prominent initiatives is the *agent protocol* proposed by the AI Engineer Foundation [30]. Several relevant projects, such as AutoGPT [1] (a popular framework for building agents that focus on task automation), are adopting similar protocols. A committee within the IEEE Standards Association (IEEE SA) is also currently working on an agent interface standard [8]. In terms of technology, while (i) Git has been established as the *de facto* technology behind collaborative development in Codeware and (ii) innovative technologies are being proposed to enable version controlling machine learning models in Neuralware [49], there is no established technology to enable collaborative development of FMware assets. Similarly, while established hubs exist for Codeware (e.g., GitHub) and Neuralware (e.g., Hugging Face), none exist for FMware assets.

**Innovation path.** Novel processes and technologies for versioning FMware assets need to be designed. Accomplishing this design requires determining what exact pieces of information need to be version-controlled and how frequently. In particular, the granularity at which prompts are defined must be examined carefully. Establishing standards and protocols for describing prompts and agents is a community endeavor that must be carried out to enable inter-organization reuse and interoperability. Once solutions to these problems are in place, a proper sharing platform (e.g., community-driven hubs) should be designed to foster the creation of a healthy ecosystem of high-quality FMware assets.

In our platform, we started approaching this challenge by enabling version control and granular role-based access control for all FMware assets. Particularly, in the case of prompts, each one of its logical components (e.g., instructions, examples, persona) is individually version- and access-controlled, enabling fine-grained asset management and reuse. In addition, all these components can be made available in our internal prompt hub, which can be accessed across an organization to enable inner-source initiatives.

### 3.7 Operation & Semantic Observability

⚙️ Serving & Ops    📊 Productivity    ⚠️ Risk    💰 Operational Cost

**Description.** FMware demands enhanced observability for a comprehensive understanding of its capability and behaviors since monitoring tools for classic software (e.g., tracking service status or execution logs) often fail to provide insights into the semantic outputs of FMware. Such enhanced observability is achieved by two approaches that we have coined as *enhanced operation telemetry* and *semantic signal telemetry*.

Enhanced operation telemetry includes monitoring for data drift, model parameters, token consumptions, groundedness, etc. during operations [78]. Such telemetry data enables analyzing the behaviors of FMware, identifying and ensuring the accuracy of the FMware, and preventing business-impacting issues when the data landscape changes, especially considering the inherently high cardinality of involved data. Furthermore, an FMware is typically inherently connected to multiple systems (e.g., legacy systems, databases, or even other FMware), and leverages the data produced by such a system as input to generate its own output, resulting in a complex system. Enhanced operation metrics in such a complex system help FMware developers navigate dynamic interactions



among FMware and environments, and thus streamlines debugging activities. However, existing observability solutions for Codeware cannot provide comprehensive observability features for the unique metrics and complex interactions of FMware [27].

Semantic signal telemetry monitors explicit and implicit feedback of users to assess the quality of FMware in production, as the quality of FMware output is difficult to measure (see Section 3.9). Explicit user feedback often takes the form of a rating system (e.g., thumb-up/thumb-down or scale ratings). However, constantly asking end-users to provide explicit feedback to the FMware output is disruptive to the user experience. Implicit user feedback often involves monitoring end-users' reactions to the FMware output, such as the modification of generated results. However, defining accurate proxy behavior to user sentiment can be difficult.

**State of practice.** Recently many efforts have been made to enhance the observability of FMware. LangSmith, a closed-source platform from the popular FMware framework LangChain's community, provides observability in complex LangChain applications through logging runs, managing nested calls, and tracing component interactions [54]. However, achieving a level of observability that can comprehensively trace, monitor, and analyze the dynamics of FMware ecosystems, especially for Agentware with multiple agents beyond LangChain's single-agent and chain mechanism, remains a formidable challenge. Specific applications such as GitHub Copilot applies semantic signal telemetry by monitoring the editing behavior of users after a piece of code is generated, as an implicit feedback signal. However, there is no generalizable solutions for semantic signal telemetry for FMware.

**Innovation path.** Enhanced metrics of diverse aspects for observability are necessary to assist developers in understanding the capabilities, analyze behaviors, and ensure the performance of FMware, given the dynamics and complexity of FMware. Our FMArts platform includes an *operations studio* that facilitates observability by supporting the monitoring, logging, and alerting of metrics that are unique to FMware. However, more efforts need to be made to advance generalizable semantic signal telemetry.

### 3.8 Performance Engineering

⚙️ Serving & Ops.    📈 Productivity    💰 Operational Cost

**Description.** Deploying and serving enterprise FMware introduces significant performance challenges, mainly due to the vast size of the underlying FMs and the need for multiple, high-performance interactions [13]. In line with recent research, we find that larger prompts or more complex FMs increase inference time and performance drops with longer prompts from multi-round conversations, which are common in enterprise FMware. Furthermore, when FMware uses externally-hosted FMs, larger prompts or multi-round conversations involve increased token usage and, as a consequence, the cost and latency of operating the FMware [47] increases. This is problematic due to token limits/context length and inference speed throttling enforced by FM providers, which impact the performance of the deployed FMware. Thus, tackling these performance issues is crucial for successfully implementing trustworthy FMware.

In addition, in an enterprise context, an FMware may involve multiple FMs and multiple rounds of inference for each FM. The

number of rounds of inference may not be predictable before execution due to the potential dynamic nature of FM-based planning and orchestration. Hence, even if the performance of each FM is optimized and meets the model level Service Level Agreements (SLAs), the FMware may still fail the application level SLA. The problem compounds in an Agentware, where each agent is expected to continuously plan, learn, and evolve, requiring unforeseeable amounts of inference from multiple FMs.

Another key concern in an enterprise setting is that when the FMs are self-hosted on an enterprise's private cloud, the computing resources (e.g., GPUs) available to a given FMware are often limited. In such a case not all needed FMs can be loaded into GPUs at the same time. Loading and off-loading FMs by demand are therefore needed and could take tens of seconds each time, posing significant performance challenges to FMware. Furthermore, FMware, which is composed of multiple agents, needs a data dependency-oriented execution plan since one agent might depend on the outcome of another agent. In such a setting, optimally loading and unloading the FMs required by different agents into the GPU, given that some FMs may have a cold start problem for a given agent (since the FM may not be finetuned) within the constraints of the FMware's SLA, poses significant performance challenges to enterprise FMware.

**State of practice.** Much of the existing research focuses on solving performance bottlenecks at a single-round inference level (i.e., FMOps level). There is little progress made in the FMware level with multiple models and multiple rounds of inference (i.e., FMware-Ops level). One strategy used to ameliorate performance issues in FMware level is selecting FMs with fewer parameters to power the FMware. For instance, in Meta, when building an FMware for code authoring, they purposefully choose the 1.3B InCoder FM over the 6.7B model to avoid inference latency despite the drop in accuracy [72]. Another method involves compressing the FMs through quantization, distillation, or pruning, reducing the FM's size to improve inference time [116]. However, how to effectively compress the FM without sacrificing its performance remains a challenge. Additionally, researchers explore prompt compression [48], aiming to shorten input prompts while preserving essential information, though such practice risks losing critical semantic details, especially in logical reasoning tasks. A promising area is semantic caching [34], which stores responses to previous prompts for quick retrieval in similar or identical queries, enhancing the efficiency and speed of enterprise FMware interactions.

**Innovation path.** Existing practices treat FM inference acceleration as a standalone endeavour and FMware is parsed into a series of imperative statements used in common programming languages (e.g., Python). Such a representation will lose the intent behind developers' reasoning about optimal task placement associated with data locality, leaving the optimization made only by compilers. To cope with this challenge, we propose that FMware should be defined and presented in a declarative manner with a higher level of abstraction. Such abstractions could capture developers' intent and expose advanced application graph-level optimization [112]. In addition, optimizing the performance at FMware level also takes into account FM characteristics, hardware heterogeneity, and data movement cost, providing a global performance view compared with accelerations of individual FM inference speed. More research should

be conducted towards these joint optimizations between the front-end intent-aware representation, the back-end application graph performance optimization, and the FM inference accelerations. Our FMArts platform represents the developed FMware in a declarative manner. In FMArts, the SOP-based workflow (either created manually or by an FM) specifies the operating procedure, which is represented as an execution graph. This execution graph is passed on to a runtime component called *fusion runtime*, which determines the optimal way to execute the SOP-based workflow through various graph and performance optimizations (e.g., elastic resource auto-scaling). These optimizations are geared towards decreasing the FM inference latency and data movement, as well as maximizing hardware utilization while still meeting SLAs. Ultimately, our fusion runtime increases the overall performance of FMware systems by efficiently using the hardware. Finally, the fusion runtime also supports a special type of cluster, which we call *uni-cluster*. A uni-cluster is a single cluster that is specially designed to serve multiple purposes, such as FM training, fine-tuning, serving, and FM-based agent customization. In our experience, uni-clusters drastically simplify the management of computing resources, resulting in increased operations (Ops) team productivity.

### 3.9 Quality under Non-determinism

🔧 Design, Dev. & Test    ⚠️ Productivity    ⚠️ Risk

**Description.** Unlike the conventional AI models in Neuralware which are often used in predictive tasks (e.g., classification, regression), FMware often utilizes FMs in tasks that are generative in nature (e.g., content creation, reasoning and planning). Such difference poses two unique challenges in evaluating the quality of FMware. First, there exists no ground truth for a given input. Predictive tasks often have a sole ground truth that can be compared against and allow quantitative evaluation of the quality of the Neuralware. On the contrary, generative tasks often have more than one “correct” answer, which can be semantically similar (e.g., for translation tasks) or drastically different (e.g., for creative writing tasks). Hence, it is challenging to define clear “pass/fail” criteria for FMware. Second, now every test needs to deal with non-determinism and becomes flaky [79, 80]. Ensuring the reproducibility of FMware is even more challenging than that of Neuralware, invalidating the premise of using test suites to evaluate the quality of FMware as it is neither practical to list all potential variants of acceptable outputs [73], nor will the testing results remain consistent across different runs. In addition, testing FMware incurs substantial costs, as every request to an FM carries a financial burden (e.g., 1-2 cents to run [79]).

**State of practice.** The current practices of evaluating the quality of outputs from FMware either rely on human efforts, or involve automated validation against a reference solution using another high-performing FM such as GPT-4 [64]. For example, LangChain documentation suggests using several FM-based evaluators and ensuring reproducibility by setting a low temperature [55]. But even setting the temperature to zero does not guarantee determinism of the evaluator output [77]. Moreover, leveraging GPT-4 as an evaluator to assess the quality of outputs generated by FMware only achieves a correlation of 0.51 with human evaluation result [64], posing risks on the reliability of the evaluation result.

**Innovation path.** Novel approaches to define and evaluate the quality of FMware, with reduced human effort yet high reliability are needed. FMArts offers experiment and testing features to evaluate multiple FMware variants (e.g., prompt, model, and inference parameter combination) against a testing set of inputs. However, it is still challenging to quantify and summarize the test results without heavy human evaluation efforts.

### 3.10 Siloed Tooling & Lack of Process

🔧 All phases (cross-cutting)    ⚠️ Productivity    ⚠️ Risk    ⚠️ Operational Cost

**Description.** The complexity of developing FMware has led to an explosion of siloed tools that only tackle a specific aspect of the lifecycle [79]. For instance, there exist individual tools to support data alignment (e.g., Snorkel AI [85]), prompt engineering (e.g., Prompt IDEs from xAI [104] and IBM [91]), and agent orchestration (e.g., Microsoft AutoGen Studio [100]).

The use of siloed tools slows down development due to cognitive overload, context switching, and the need to implement integration (glue) code for several interfaces that are managed by different offering providers. Consequently, the pace at which software solutions can produce business value is also slowed down. In our experience, the siloed tool use also hinders the trustworthiness of FMware applications, as the lack of standardized development tools and practices leads to inefficient compliance & governance.

**State of practice.** No lifecycle engineering solution currently exists for developing FMware [79]. We do acknowledge, however, the ongoing efforts of key industry players (e.g., PromptFlow by Microsoft [70] and LangSmith [45] by Langchain). In the research space, we also observe a few initiatives, such as (Bee)\* [68] and Prompt Sapper [22]. However, current solutions either cannot cover the full lifecycle of FMware, missing critical considerations such as compliance and trustworthiness, or offer little extensibility and flexibility to integrate richer functionalities from the community.

**Innovation path.** Scalable and efficient development of FMware can only be accomplished by means of a truly unified platform that provides cradle-to-grave lifecycle support and offers flexibility for choosing between capabilities and integrating them. Our platform FMArts represents our effort towards achieving this goal.

## 4 CONCLUSION

The goal of this paper is to raise awareness of SE4FMware challenges that are observed in an industrial setting. But where should we go from here? Fully tackling those challenges requires the design of breakthrough technologies, techniques, and standards. We hope that our ten challenges will foster further discussions and knowledge sharing between software engineering researchers in industry and academia. In particular, we believe that certain challenges such as *limited collaboration support* (Section 3.6) can only be overcome by creating community-wide initiatives (e.g., a strategic consortium and interest groups) that carry out open discussions and prescribe the involvement of the broader community of software developers, industries, and open-source. More generally, we hope that the disclosure of our challenges, progress, and solutions will foster further explorations and discussions around the many challenges of this nascent and yet important domain of FMware.

## REFERENCES

- [1] [n. d.]. AutoGPT Documentation. <https://docs.agpt.co/>. Accessed 01-17-2024.
- [2] [n. d.]. BabyAGI. <https://github.com/yoheinakajima/babyagi>. Accessed 17-01-2024.
- [3] [n. d.]. CrewAI: Framework for orchestrating role-playing, autonomous AI agents. <https://github.com/joaoimdmoura/crewai>. Accessed 17-01-2024.
- [4] [n. d.]. MiniAGI - a minimal general-purpose autonomous agent based on GPT-3.5 / GPT-4. <https://github.com/muellerberndt/mini-agi>. Accessed 17-01-2024.
- [5] A. Alford. 2024. OpenAI Releases New Embedding Models and Improved GPT-4 Turbo. <https://www.infoq.com/news/2024/02/openai-model-updates/>.
- [6] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, and T. Zimmermann. 2019. Software engineering for machine learning: a case study. In *Proceedings of the 41st International Conference on Software Engineering (Montreal, Canada) (ICSE-SEIP '19)*. IEEE Press, 291–300.
- [7] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, et al. 2016. Concrete Problems in AI Safety. arXiv:1606.06565 [cs.AI]
- [8] Artificial Intelligence Standards Committee (C/AISC). [n. d.]. P3394: Standard for Large Language Model Agent Interface. <https://standards.ieee.org/ieee/3394/11377>. Accessed 02-10-2024.
- [9] S.H Bach, D. Rodriguez, Y. Liu, C. Luo, H. Shao, et al. 2019. Snorkel drybell: A case study in deploying weak supervision at industrial scale. In *Proc. of Int. Conf. on Management of Data (SIGMOD)*. 362–375.
- [10] Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, et al. 2022. Constitutional AI: Harmlessness from AI Feedback. arXiv:2212.08073 [cs.CL]
- [11] M. Benjamin, P. Gagnon, N. Rostamzadeh, C. Pal, Y. Bengio, and A. Shee. 2019. Towards standardization of data licenses: The montreal data license. arXiv preprint arXiv:1903.12262 (2019).
- [12] S. Biegel, R. El-Khatib, L. O. V. B. Oliveira, M. Baak, and N. Aben. 2021. Active weasul: improving weak supervision with active learning. arXiv preprint arXiv:2104.14847 (2021).
- [13] R. Bommasani, D.A Hudson, E. Adeli, Altman, et al. 2021. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021).
- [14] H.B. Braiek and F. Khomh. 2020. On testing machine learning programs. *Journal of Systems and Software* 164 (2020), 110542.
- [15] D. Brajovic, N. Renner, V.P. Goebels, P. Wagner, B. Fresz, et al. 2023. Model Reporting for Certifiable AI: A Proposal from Merging EU Regulation into AI Development. arXiv preprint arXiv:2307.11525 (2023).
- [16] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [17] M. Casey. 2023. LLMs high priority for Enterprise Data Science, but concerns remain. <https://snorkel.ai/poll-data-llm-high-priority-enterprise-data-science-concerns-remain>. Accessed 01-29-2024.
- [18] J. Chakraborty, S. Majumder, and T. Menzies. 2021. Bias in machine learning software: why? how? what to do?. In *Proc. of ACM Joint Meeting on ESEC/FSE*. ACM, 429–440.
- [19] K.K Chang, M. Cramer, S. Soni, and D. Bamman. 2023. Speak, memory: An archaeology of books known to chatgpt/gpt-4. arXiv preprint arXiv:2305.00118 (2023).
- [20] J. Chen, N. Yoshida, and H. Takada. 2023. An investigation of licensing of datasets for machine learning based on the GQM model. arXiv preprint arXiv:2303.13735 (2023).
- [21] L. Chen, M. Zaharia, and J. Zou. 2023. How is ChatGPT's behavior changing over time? arXiv:2307.09009 [cs.CL]
- [22] Y. Cheng, J. Chen, Q. Huang, Z. Xing, X. Xu, et al. 2023. Prompt Sapper: A LLM-Empowered Production Tool for Building AI Chains. *ACM Trans. on Softw. Eng. and Methodology* (Dec. 2023).
- [23] T. Claburn. 2023. GitHub, Microsoft, OpenAI fail to wriggle out of Copilot copyright lawsuit. [https://www.theregister.com/2023/05/12/github\\_microsoft\\_openai\\_copilot/](https://www.theregister.com/2023/05/12/github_microsoft_openai_copilot/). Accessed 01-17-2024.
- [24] D. Contractor, D. McDuff, J.K. Haines, J. Lee, C. Hines, et al. 2022. Behavioral use licensing for responsible AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 778–788.
- [25] B. Deiseroth, M. Deb, S. Weinbach, M. Brack, P. Schramowski, et al. 2023. AtMan: Understanding Transformer Predictions Through Memory Efficient Attention Manipulation. arXiv:2301.08110 [cs.LG]
- [26] M. Deng, J. Wang, C.-P. Hsieh, Y. Wang, H. Guo, et al. 2022. RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning. In *Proc. of Conf. on EMNLP*. ACL, 3369–3391.
- [27] Dynatrace. 2023. What is Observability? Not just logs, metrics and traces. <https://www.dynatrace.com/news/blog/what-is-observability-2>. Accessed: 01-31-2024.
- [28] L. Edwards. 2021. The EU AI Act: a summary of its significance and scope. *AI (the EU AI Act)* 1 (2021).
- [29] C. Fernando, D. Banarse, H. Michalewski, S. Osindero, and T. Rocktäschel. 2023. Promptbreeder: Self-Referential Self-Improvement Via Prompt Evolution. arXiv:2309.16797 [cs.CL]
- [30] AI Engineer Foundation. 2023. Agent Protocol. <https://agentprotocol.ai>. Accessed: 01-31-2024.
- [31] Linux Foundation. [n. d.]. SPDX Announces 3.0 Release Candidate with New Use Cases. <https://www.linuxfoundation.org/press/spdx-sbom-3-release-candidate>. Accessed 02-07-2024.
- [32] Martin Fowler. 2002. *Patterns of Enterprise Application Architecture*. Addison-Wesley Longman Publishing Co., Inc., USA.
- [33] T. Gao, A. Fisch, and D. Chen. 2021. Making Pre-trained Language Models Better Few-shot Learners. In *Proc. of Annual Meeting of the ACL-IJCNLP (ACL)*, 3816–3830.
- [34] I. Gim, G. Chen, S.S. Lee, N. Sarda, A. Khandelwal, et al. 2023. Prompt Cache: Modular Attention Reuse for Low-Latency Inference. arXiv preprint arXiv:2311.04934 (2023).
- [35] Github. [n. d.]. GitHub Next | Copilot Workspace. <https://githubnext.com/projects/copilot-workspace/>. Accessed 02-06-2024.
- [36] Google. [n. d.]. Permissions on Android. <https://developer.android.com/guide/topics/permissions/overview>. Accessed 02-06-2024.
- [37] Grand View Research. 2023. Large Language Model Market Size, Share & Trends Analysis Report By Application (Customer Service, Content Generation), By Deployment, By Industry Vertical, By Region, And Segment Forecasts, 2024 - 2030. <https://www.grandviewresearch.com/industry-analysis/large-language-model-llm-market-report>.
- [38] R. Grosse, J. Bae, C. Anil, N. Elhage, A. Tamkin, et al. 2023. Studying large language model generalization with influence functions. arXiv preprint arXiv:2308.03296 (2023).
- [39] S. Gunasekar, Y. Zhang, J. Aneja, C.C.T. Mendes, A. Del Giorno, et al. 2023. Textbooks Are All You Need. arXiv preprint arXiv:2306.11644 (2023).
- [40] Ł. Górski and S. Ramakrishna. 2023. Challenges in Adapting LLMs for Transparency: Complying with Art. 14 EU AI Act. (2023).
- [41] A.E. Hassan, B. Adams, F. Khomh, N. Nagappan, and T. Zimmermann. 2023. FM+SE Vision 2030. <https://fmse.io/>. Accessed 01-17-2024.
- [42] S. Hong, M. Zhuge, J. Chen, X. Zheng, Y. Cheng, et al. 2023. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. arXiv:2308.00352 [cs.AI]
- [43] O. Honovich, U. Shaham, S.R. Bowman, and O. Levy. 2022. Instruction Induction: From Few Examples to Natural Language Task Descriptions. arXiv:2205.10782 [cs.CL]
- [44] C.-J. Hsieh, S. Si, F.X. Yu, and I.S. Dhillon. 2023. Automatic Engineering of Long Prompts. arXiv:2311.10117 [cs.AI]
- [45] LangChain Inc. [n. d.]. LangSmith. <https://www.langchain.com/langsmith>. Accessed 01-18-2024.
- [46] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, et al. [n. d.]. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* 55, 12, Article 248 (n. d.).
- [47] A.Q Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D.S. Chaplot, et al. 2023. Mistral 7B. arXiv preprint arXiv:2310.06825 (2023).
- [48] H. Jiang, Q. Wu, C.Y. Lin, Y. Yang, and L. Qiu. 2023. Lmlingua: Compressing prompts for accelerated inference of large language models. arXiv preprint arXiv:2310.05736 (2023).
- [49] N. Kandpal, B. Lester, M. Muqeeth, A. Mascarenhas, M. Evans, et al. 2023. Git-Theta: a git extension for collaborative development of machine learning models. In *Proc. of Int. Conf. on Machine Learning*. Article 642, 12 pages.
- [50] O. Khattab, A. Singhvi, P. Maheshwari, Z. Zhang, K. Santhanam, et al. 2023. DSPy: Compiling Declarative Language Model Calls into Self-Improving Pipelines. arXiv preprint arXiv:2310.03714 (2023).
- [51] F. Khomh, H. Li, M. Lamothe, M.A. Hamdaqa, J. Cheng, et al. 2023. Software Engineering for Machine Learning Applications (SEMLA) 2023. <https://semmla.polyml.ca/>. Accessed 01-17-2024.
- [52] S. Krishnan, M.J. Franklin, K. Goldberg, and E. Wu. 2017. BoostClean: Automated Error Detection and Repair for Machine Learning. arXiv:1711.01299 [cs.DB]
- [53] S. Krishnan, J. Wang, E. Wu, M.J. Franklin, and K. Goldberg. 2016. ActiveClean: Interactive Data Cleaning While Learning Convex Loss Models. arXiv:1601.03797 [cs.DB]
- [54] LangChain, Inc. 2024. LangSmith Documentation. <https://docs.smith.langchain.com/>. Accessed: 01-31-2024.
- [55] LangChain, Inc. 2024. Off-the-shelf LangChain Evaluators. <https://docs.smith.langchain.com/evaluation/evaluator-implementations>. Accessed: 01-31-2024.
- [56] Langchain Team. 2023. LANGCHAIN BLOG: OpenAI's Bet on a Cognitive Architecture. <https://blog.langchain.dev/openais-bet-on-a-cognitive-architecture/>. Accessed: 2024-02-08.
- [57] Y. LeCun. 2022. What is the future of AI. [https://www.facebook.com/watch/live/?ref=watch\\_permalink&v=2219848494820560](https://www.facebook.com/watch/live/?ref=watch_permalink&v=2219848494820560). Accessed: 2024-02-08.
- [58] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, et al. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 9459–9474.
- [59] Y. Li, L. Meng, L. Chen, L. Yu, D. Wu, et al. 2022. Training Data Debugging for the Fairness of Machine Learning Software. In *Proc. of Int. Conf. on Softw. Eng. (ICSE)*. 2215–2227.
- [60] Y. Liang, C. Wu, T. Song, W. Wu, Y. Xia, et al. 2023. TaskMatrix.AI: Completing Tasks by Connecting Foundation Models with Millions of APIs.

- arXiv:2303.16434 [cs.AI]
- [61] N.F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, et al. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172* (2023).
- [62] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig. 2023. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Comput. Surv.* 55, 9, Article 195 (2023).
- [63] Y. Liu. 2023. The importance of human-labeled data in the era of LLMs. *arXiv preprint arXiv:2306.14910* (2023).
- [64] Y. Liu, D. Iter, Yi. Xu, S. Wang, R. Xu, and C. Zhu. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634* (2023).
- [65] Llamaindex. [n. d.]. LlamaIndex (formerly GPT Index) is a data framework for your LLM applications. [https://github.com/run-llama/llama\\_index](https://github.com/run-llama/llama_index). Accessed 02-07-2024.
- [66] S. Longpre, R. Mahari, A. Chen, N. Obeng-Marnu, D. Sileo, et al. 2023. The Data Provenance Initiative: A Large Scale Audit of Dataset Licensing & Attribution in AI. *arXiv preprint arXiv:2310.16787* (2023).
- [67] S. Maekawa, D. Zhang, H. Kim, S. Rahman, and E. Hruschka. 2022. Low-resource interactive active labeling for fine-tuning language models. In *Findings of the Assoc. for Computational Linguistics: EMNLP 2022*. 3230–3242.
- [68] L. Martie, J. Rosenberg, V. Demers, G. Zhang, O. Bhardwaj, et al. 2023. Rapid Development of Compositional AI. In *Proc. of Int. Conf. Softw. Eng.: New Ideas and Emerging Results (ICSE-NIER)*.
- [69] Microsoft. [n. d.]. microsoft/semantic-kernel: Integrate cutting-edge LLM technology quickly and easily into your apps. <https://github.com/microsoft/semantic-kernel>. Accessed 01-30-2024.
- [70] Microsoft. 2023. Prompt flow documentation. <https://microsoft.github.io/promptflow>. Accessed 02-06-2024.
- [71] M. Mozes, X. He, B. Kleinberg, and L.D. Griffin. 2023. Use of llms for illicit purposes: Threats, prevention measures, and vulnerabilities. *arXiv preprint arXiv:2308.12833* (2023).
- [72] V. Murali, C. Maddila, I. Ahmad, M. Bolin, D. Cheng, et al. 2023. CodeCompose: A Large-Scale Industrial Deployment of AI-assisted Code Authoring. *arXiv preprint arXiv:2305.12050* (2023).
- [73] A. Nguyen-Duc, B. Cabrero-Daniel, A. Przybylek, C. Arora, D. Khanna, et al. 2023. Generative Artificial Intelligence for Software Engineering—A Research Agenda. *arXiv preprint arXiv:2310.18648* (2023).
- [74] NIST. 2021. Improving the Nation’s Cybersecurity: NIST’s Responsibilities Under the May 2021 Executive Order. <https://www.nist.gov/itl/executive-order-improving-nations-cybersecurity/software-supply-chain-security-guidance-0>. Accessed 02-01-2024.
- [75] C. Novelli, F. Casolari, A. Rotolo, M. Taddeo, and L. Floridi. 2023. Taking AI risks seriously: a new assessment model for the AI Act. *AI & SOCIETY* (2023).
- [76] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [77] S. Ouyang, J.M. Zhang, M. Harman, and M. Wang. 2023. LLM is Like a Box of Chocolates: the Non-determinism of ChatGPT in Code Generation. *arXiv preprint arXiv:2308.02828* (2023).
- [78] V. Pamula. 2023. An Introduction to LLMops: Operationalizing and Managing Large Language Models using Azure ML. <https://techcommunity.microsoft.com/t5/ai-machine-learning-blog/an-introduction-to-llmops-operationalizing-and-managing-large/ba-p/3910996>. Accessed: 01-31-2024.
- [79] C. Parnin, G. Soares, R. Pandita, S. Gulwani, J. Rich, et al. 2023. Building Your Own Product Copilot: Challenges, Opportunities, and Needs. arXiv:2312.14231 [cs.SE]
- [80] O. Parry, G.M. Kapfhammer, M. Hilton, and P. McMinn. 2022. Surveying the developer experience of flaky tests. In *Proc. of Int. Conf. on Softw. Eng.: Software Engineering in Practice*. 253–262.
- [81] Jenny Preece, Yvonne Rogers, and Helen Sharp. 2002. *Interaction Design* (1st ed.). John Wiley & Sons, Inc., USA.
- [82] Pallets Project. [n. d.]. Jinja. <https://palletsprojects.com/p/jinja/>. Accessed 02-06-2024.
- [83] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, et al. 2023. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. arXiv:2307.16789 [cs.AI]
- [84] G.K. Rajbahadur, E. Tuck, L. Zi, D. Lin, B. Chen, et al. 2021. Can I use this publicly available dataset to build commercial AI software?—A Case Study on Publicly Available Image Datasets. *arXiv preprint arXiv:2111.02374* (2021).
- [85] A. Ratner, S.H. Bach, H. Ehrenberg, J. Fries, S. Wu, et al. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proc. of Int. Conf. on Very Large Data Bases (VLDB Endowment)*.
- [86] J. Robinson and Andrei Voronkov. 2001. *Handbook of Automated Reasoning: Volume 1*. MIT Press, Cambridge, MA, USA.
- [87] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, et al. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv:2302.04761 [cs.CL]
- [88] S. Schoch, R. Mishra, and Y. Ji. 2023. Data Selection for Fine-tuning Large Language Models Using Transferred Shapley Values. *arXiv preprint arXiv:2306.10165* (2023).
- [89] W. Shi, A. Ajith, M. Xia, Y. Huang, D. Liu, et al. 2023. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789* (2023).
- [90] I. Shumailov, Z. Shumaylov, Y. Zhao, Y. Gal, N. Papernot, et al. 2023. The curse of recursion: training on generated data makes models forget. *Arxiv. Preprint posted online 27* (2023).
- [91] H. Strobelt, A. Webson, V. Sanh, B. Hoover, J. Beyer, et al. 2023. Interactive and Visual Prompt Engineering for Ad-hoc Task Adaptation with Large Language Models. *IEEE Trans. on Visualization and Computer Graphics* 29, 1 (2023), 1146–1156.
- [92] S. Tilga. [n. d.]. LLMs & humans: The perfect duo for data labeling. <https://toloka.ai/blog/llms-and-humans-for-data-labeling/>. Accessed 01-29-2024.
- [93] K. Valmeekam, A. Olmo, S. Sreedharan, and S. Kambhampati. 2022. Large Language Models Still Can’t Plan (A Benchmark for LLMs on Planning and Reasoning about Change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- [94] L. Wang, X. Zhang, H. Su, and J. Zhu. 2023. A Comprehensive Survey of Continual Learning: Theory, Method and Application. arXiv:2302.00487 [cs.LG]
- [95] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N.A. Smith, et al. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. arXiv:2212.10560 [cs.CL]
- [96] J. Wei, X. Wang, D. Schuurmans, M. Bosma, E. Chi, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* (2022).
- [97] W.E. Wong, V. Debroy, A. Surampudi, H. Kim, and M.F. Siok. 2010. Recent Catastrophic Accidents: Investigating How Software was Responsible. In *Int. Conf. on Secure Softw. Integration and Reliability Improvement*. 14–22.
- [98] FOSSology Workgroup. 2023. FOSSology. <https://www.fossology.org/>. Accessed 02-07-2024.
- [99] OpenDataology workgroup. 2022. OpenDataology. <https://github.com/OpenDataology/>. Accessed 02-07-2024.
- [100] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, et al. 2023. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. arXiv:2308.08155 [cs.AI]
- [101] T. Wu, E. Jiang, A. Donsbach, J. Gray, A. Molina, et al. 2022. PromptChainer: Chaining Large Language Model Prompts through Visual Programming. In *Extended Abstracts of Conf. on Human Factors in Computing Systems*. ACM.
- [102] T. Wu, M. Terry, and C.J. Cai. 2022. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. In *Proc. of Conf. on Human Factors in Computing Systems (CHI)*. ACM.
- [103] Z. Wu, L. Qiu, A. Ross, E. Akyurek, B. Chen, et al. 2023. Reasoning or Reciting? Exploring the Capabilities and Limitations of Language Models Through Counterfactual Tasks. arXiv:2307.02477 [cs.CL]
- [104] xAI Team. 2023. PromptIDE. <https://x.ai/prompt-ide/>. Accessed 01-18-2024.
- [105] Jiawen Xiong, Yong Shi, Boyuan Chen, Filipe R. Cogo, and Zhen Ming (Jack) Jiang. 2022. Towards build verifiability for Java-based systems. In *Proceedings of the 44th International Conference on Software Engineering: Software Engineering in Practice (Pittsburgh, Pennsylvania) (ICSE-SEIP ’22)*. Association for Computing Machinery, New York, NY, USA, 297–306. <https://doi.org/10.1145/3510457.3513050>
- [106] H. Xu, Y. Chen, Y. Du, N. Shao, W. Yanggang, et al. 2022. GPS: Genetic Prompt Search for Efficient Few-Shot Learning. In *Proc. of Conf. on EMNLP. ACL*, 8162–8171.
- [107] Q. Xu, F. Hong, B. Li, C. Hu, Z. Chen, et al. 2023. On the Tool Manipulation Capability of Open-source Large Language Models. arXiv:2305.16504 [cs.CL]
- [108] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629* (2022).
- [109] Q. Ye, M. Axmed, R. Pryzant, and F. Khani. 2023. Prompt Engineering a Prompt Engineer. arXiv:2311.05661 [cs.CL]
- [110] Y. Yu, L. Kong, J. Zhang, R. Zhang, and C. Zhang. 2022. AcTune: Uncertainty-based active self-training for active fine-tuning of pretrained language models. In *Proc. of Conf. of the North American Chapter of the Assoc. for Computational Linguistics: Human Language Technologies*. 1422–1436.
- [111] S. Zhang, L. Dong, X. Li, S. Zhang, X. Sun, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792* (2023).
- [112] L. Zheng, L. Yin, Z. Xie, J. Huang, C. Sun, et al. 2023. Efficiently Programming Large Language Models using SGLang. *CoRR* (2023).
- [113] C. Zhou, P. Liu, P. Xu, S. Iyer, J. Sun, et al. 2023. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206* (2023).
- [114] W. Zhou, Y.E. Jiang, L. Li, J. Wu, T. Wang, et al. 2023. Agents: An Open-source Framework for Autonomous Language Agents. arXiv:2309.07870 [cs.CL]
- [115] Y. Zhou, A.I. Muresanu, Z. Han, K. Paster, S. Pitis, et al. 2023. Large Language Models Are Human-Level Prompt Engineers. arXiv:2211.01910 [cs.LG]
- [116] X. Zhu, J. Li, Y. Liu, C. Ma, and W. Wang. 2023. A survey on model compression for large language models. *arXiv preprint arXiv:2308.07633* (2023).

Received 2024-02-08; accepted 2024-04-18