

A Meta-Summary of Challenges in Building Products with ML Components – Collecting Experiences from 4758+ Practitioners

Nadia Nahar
nadian@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, PA, USA

Haoran Zhang
Carnegie Mellon University
Pittsburgh, PA, USA

Grace Lewis
Carnegie Mellon Software
Engineering Institute
Pittsburgh, PA, USA

Shurui Zhou
University of Toronto
Toronto, Ontario, Canada

Christian Kästner
Carnegie Mellon University
Pittsburgh, PA, USA

ABSTRACT

Incorporating machine learning (ML) components into software products raises new software-engineering challenges and exacerbates existing challenges. Many researchers have invested significant effort in understanding the challenges of industry practitioners working on building products with ML components, through interviews and surveys with practitioners. With the intention to aggregate and present their collective findings, we conduct a meta-summary study: We collect 50 relevant papers that together interacted with over 4758 practitioners using guidelines for systematic literature reviews. We then collected, grouped, and organized the over 500 mentions of challenges within those papers. We highlight the most commonly reported challenges and hope this meta-summary will be a useful resource for the research community to prioritize research and education in this field.

1 INTRODUCTION

After decades of effort in machine learning (ML) research to build better models, researchers from industry and academia have recently started to shift their attention to improving how to build software products with such models. Incorporating a ML component into a software product is often argued to be harder than incorporating traditional functional components, because of the specific characteristics of machine learning (e.g., based on data, no specifications in the traditional sense, fairness concerns) and how they impact the entire life cycle of the product [42, 50, 59, 84, 112]. While the traditional software development process has challenges of its own, bringing ML into the picture is argued to break a lot of existing software architecture and engineering assumptions [50, 59]. This leads initiatives to rethink existing processes and practices and shift priorities in software teams. As a result, we keep hearing from practitioners on how they perceive building, deploying, and incorporating machine learning in software products as a challenge, even when the initial ML research and model prototypes seemed promising.

While some practitioners give talks on challenges or write experience papers (e.g., examples in academic venues surveyed elsewhere [67, 85]), researchers have also been actively studying the challenges faced by practitioners when building software products with ML components across many projects. In recent years, many researchers have *interviewed* or *surveyed* practitioners to identify what has really changed for them with the introduction of machine learning, often with the goal of identifying challenges, research

Requirements Engineering: Lack of AI literacy causes unrealistic expectations from customers, managers, and even other team members • Vagueness in ML problem specifications makes it difficult to map business goals to performance metrics • Regulatory constraints specific to data and ML introduce additional requirements that restrict development

Architecture, Design, and Implementation: Transitioning from a model-centric to a pipeline-driven or system-wide view is considered important for moving into production, but a difficult paradigm shift for many teams • ML adds substantial design complexity with many, often implicit, data and tooling dependencies, and entanglements due to a lack of modularity • Difficulty in scaling model training and deployment on diverse hardware • While monitorability and planning for change are often considered important, they are mostly considered only late after launching

Model Development: Model development benefits from engineering infrastructure and tooling but provided infrastructure and technical support are limited in many teams • Code quality is not standardized in model development tools, leading to conflicts about code quality

Data Engineering: Data quality is considered important, but difficult for practitioners and not well supported by tools • Internal data security and privacy policies restrict data access and use • Although training-serving skew is common, many teams lack support for its required detection and monitoring • Data versioning and provenance tracking are often seen as elusive, with not enough tool support

Quality Assurance: Testing and debugging ML models is difficult due to lack of specifications • Testing of model interactions, pipelines, and the entire system is considered challenging and often neglected • Testing and monitoring models in production are considered important but difficult, and often not done • There are no standard processes or guidelines on how to assess system qualities such as fairness, security, and safety in practice

Process: Development of products with ML component(s) is often ad-hoc, lacking well-defined processes • The uncertainty in ML development makes it hard to plan and estimate effort and time

Organization and Teams: Building products with ML components requires diverse skill sets, which is often missing in development teams • Many teams are not well prepared for the extensive interdisciplinary collaboration and communication needed in ML products • ML development can be costly and resource limits can substantially curb/limit efforts • Lack of organizational incentives, resources, and education hampers achieving all system-level qualities

Table 1: Overview of Identified Challenges

opportunities, and best practices in a rapidly changing field. While some studies focus on specific aspects, such as challenges regarding




Figure 1: Research Method

architecture [110], collaboration [78], or fairness [37, 95], many others explore challenges more broadly. Many of these studies have identified similar challenges. We believe that we have reached a point where practices have settled and research on challenges approaches saturation – we think that now is a good time to step back and survey the collective findings of the research community.

In this paper, we aim to consolidate knowledge about challenges in the practice of building software products with ML components, with a systematic literature survey of existing studies that *interviewed* or *surveyed* industry practitioners across multiple projects. We identified 50 studies of which, 30 conducted interviews, 11 conducted surveys, and nine did both, with a total of over 4758 identified participants (seven studies did not report the number of participants; some participants may have participated in multiple studies). Using the *meta-summary research method* [25, 96, 104], we analyze, organize, and synthesize findings across all these studies (as shown in Figure 1), answering the overall research question: **What are the challenges experienced by industry practitioners in building software products with ML components?**

We group the challenges found in the meta-summary into categories. In a nutshell, we find practitioners struggle in different product development stages: (1) requirements engineering, (2) architecture, design, and implementation, and (3) quality assurance. We also find several engineering challenges in ML-specific stages, in particular (4) model development, and (5) data engineering. Other issues relate to cross-cutting concerns related to (6) process, and (7) organization and teams. Following the meta-summary method, we present and organize the challenges mentioned by the practitioners in the original papers as they have been reported, without attempting to speculate or pass our own judgments on the findings. Table ?? contains a summary of the findings. We conclude the paper with a brief discussion, reflecting our own views.

Our key contribution is the meta-summary, which we present in narrative form in this paper. We additionally provide details with clear traceability to findings and papers as supplementary documents [77].

2 SCOPING AND RELATED WORK

With the advance of ML techniques, many organizations have invested substantial efforts in building products with ML components. While there is a large amount of research that focuses entirely on the challenges that data scientists face in their model-development work (e.g., development responsibilities [52], data exploration [63, 75], data-science processes [30, 70], development in notebooks [20, 33, 88], AutoML [122]), another body of work focuses on the challenges of building products with those models, often with interdisciplinary teams, and placing substantial attention on qualities like safety and observability. The latter work, which forms the scope of this survey, moves beyond the model-centric view of classic data-science workflows and considers building automated pipelines and entire software systems with many ML and non-ML components, as well as the engineering challenges involved. It emerges in a growing research community often named *Software Engineering for Machine Learning (SE4ML)* studying the engineering challenges of building both ML components and products that contain ML components.

Understanding practitioner needs. Academic research is often criticized for being far removed from the needs faced by practitioners in industry [9, 95, 121]. If researchers want to achieve rapid impact in industry, they need to understand what problems are important to practitioners; conversely, practitioners may attempt to attract researchers to work on their problems. Attempts to close the gap between academia and practice typically need to navigate a tradeoff between (a) investigating one or few teams in depth with findings that may not generalize or (b) exploring common problems across many teams with more shallow engagements. Focused on individual teams, we see a few ethnographic studies [86, 87], many direct collaborations with an industrial partner [53, 94], and many experience reports published by practitioners in papers [10, 32, 49, 60], talks [27, 31, 71], or blog posts [35, 114, 132]. To understand problems across teams, many researchers conduct interviews across multiple teams and organizations, e.g., [59, 66, 78, 110, 120], to be either addressed by the same researchers or reported as open problems to the community. Other researchers have focused on surveying practitioners at scale across companies and regions, e.g., [45, 58, 123, 131].

In this paper, we go one step further in aggregating and analyzing results from prior interviews and surveys with over 4758 practitioners, which we hope will help guide future research and educational activities toward challenges relevant to practitioners.

Previous literature reviews. There have been several prior literature reviews on topics related to building products with ML components. Most surveys review academic papers proposing solutions in subfields, such as testing ML components [6, 16, 39, 98, 130], safety and security [14, 39, 64], data management [91], and even trying to cover published research on SE4ML broadly [28, 69]. The closest to our work are two literature surveys that analyze practitioner experience reports published at academic conferences (not including grey literature) collecting the self-reported challenges of a few dozen teams [67, 85]. In this work, we specifically perform a meta-summary of academic papers reporting on interviews and surveys with practitioners.

3 RESEARCH METHOD

The goal of this paper is to summarize challenges in building products with ML components, accumulated from industry practitioners in prior research. To achieve this goal and answer our research question, we first define the appropriate search strategy and study selection criteria to find the relevant literature that identifies challenges through communicating with industry practitioners. We follow the guidelines for systematic literature review (SLR) for this paper collection step [51]. Then, we extract the data from the selected papers and analyze the data to complete the synthesis process. Several approaches have been explored for synthesizing qualitative research in software engineering, such as thematic synthesis, meta-ethnography, and meta-summary [40]. As our aim is to discover patterns or themes of challenges in building products with ML components, as well as get a sense of the priority of the challenges based on the frequency of reports by industry practitioners, the meta-summary method is best suited for this research problem [40, 96, 104]. The meta-summary method provides a well-balanced synthesis mechanism, which is deeper than mapping studies, and not as exhaustive as meta-ethnography, which requires significant expertise and experience with the methodology and its philosophical stance [96]. Thus, we apply the meta-summary method [25, 96, 104] to perform the quantitative aggregation of the qualitative evidence that we present as findings. Figure 1 shows the overview of the research process followed in this study.

3.1 Paper Selection

To increase reliability, reproducibility and objectivity of the process for paper selection, we follow the established procedure of conducting systematic literature reviews [51].

Relevant years. Much of the research on engineering products with ML components was inspired by the seminal 2015 ML technical debt paper by Sculley et al. [105], which outlined various engineering challenges in building and operating ML infrastructure. For completeness, we selected the year range of the papers to be from 2010 to 2022.

Publication venues. To search for papers, we select digital libraries and databases commonly used by software engineering review papers, e.g., [21, 46, 65]. We do not filter by the venue, as we expect to find papers that are published in different communities including software engineering, human-computer interaction, and machine learning. Since we aim to aggregate results from robust empirical studies, we did not include gray literature, such as blog posts, which typically reflect opinions or individual experience only. However, we did include arXiv as a data source, as it contains many relevant academic papers in this field, even if some have not been peer reviewed. Specifically, we use the following 8 data sources: IEEE Xplore (ieeexplore.ieee.org), ACM Digital library (portal.acm.org/dl.cfm), Wiley InterScience (www.interscience.wiley.com), Elsevier Science Direct (www.sciencedirect.com), SpringerLink (www.springerlink.com), EI Compendex (www.engineeringvillage.com), and arXiv (<https://arxiv.org>).

Search query. Defining the right scope and corresponding search query required some iteration. We started by assembling an initial set of 21 papers as a seed set (a common practice [26, 67]). The seed

Table 2: Paper Selection

Data Source	Initial Search Result	After Filtering by Title/Abstract and Snowballing	Final Selection
IEEE	69	30	19
ACM	48	11	10
Wiley	6	0	0
ScienceDirect	32	5	3
Engineer Village	101	3	0
Springer	6*	3	2
arXiv	79	8	5
Snowballing	-	26	11
Total	341	86	50

*abstract filtering from 5612 papers retrieved with fulltext search

set was composed of papers that we knew well from our past work in this field. We then analyzed the seed set to define the keywords needed to retrieve those and similar papers.

We realized that our research question has three aspects, and therefore to retrieve the papers that would satisfy our research question, we focused on those three parts to formulate the search query: (A) **The paper needs to mention an ML-related keyword**, since we focus on challenges introduced by ML components. (B) **The paper needs to mention a software engineering or ML deployment-related keyword**, since we focus on engineering challenges that go beyond local concerns of data scientists; for example the paper should discuss concerns related to actual product development where models are deployed and incorporated into larger software systems. Finally, (C) **the paper needs to mention surveys or interviews**, since we are interested in the challenges mentioned by industry practitioners and these are the most common relevant research methods; we are not interested in a single-team case study or ethnographic study, as the challenges found in such papers may be specific to individual products.

After adding some semantically similar terminologies, we developed the following search query fragments – A: “machine learning” OR “artificial intelligence” OR “deep learning” OR “ML component” OR “data science”; B: “software engineering” OR “software systems” OR “production-ready systems” OR “ML systems” OR “deploying ML” OR “ML deployment”; C: “interview” OR “survey” OR “questionnaire”. The final query was of the following format “A AND B AND C.”

We searched with this query within the abstract of the papers in all the digital data sources except SpringerLink, as it did not have the option to search within abstracts. For SpringerLink, we retrieved 5612 papers based on a full-text search, and subsequently used a custom script to search within the abstracts of these papers. This provided us with a total of 341 papers from all the sources (see Table 2).

This search query retrieved 18 of the 21 seed papers. Two papers were missed because the conducted interviews were not mentioned in the abstract (the abstract framed the research as a case study), and one paper was not listed within the libraries searched (only available on TechRxiv). To account for this difference we performed one round of snowballing, as explained later in this section.

Table 3: Inclusion and Exclusion Criteria


Inclusion Criteria	
I1:	Paper includes software engineering challenges for ML systems
I2:	Paper uses interview or survey with industry practitioners (software engineers, data scientists, etc.) to identify the challenges
I3:	Paper appears in a refereed publication (including conference proceedings, journal, etc.) or uploaded in arxiv in a publication format
I4:	Paper is written in English
Exclusion Criteria	
E1:	Paper has a strict ML model view and does not consider the system or product using the model
E2:	Paper interviews/surveys only non-technical people (end-users, domain experts, etc.)
E3:	Paper focuses on ML for software engineering instead of software engineering for ML systems
E4:	Paper falls in the category of gray literature: blog post, technical report, government report, webinar, poster session, presentation, etc.

Selection criteria. The initial search returned many papers that were not directly relevant to our research question. Next, we selected 86 relevant papers by reading the title and abstract, evaluating them against the inclusion and exclusion criteria (see Table 3), which we incrementally refined. Finally, we read the full paper, and once again evaluated each against the inclusion and exclusion criteria, which narrowed our set down to 39 papers. Multiple researchers participated in this process and discussed papers at the boundary.

Most of the papers that were discarded in this round were either literature surveys in the domain of *machine learning for software engineering* (i.e., using ML techniques to facilitate software engineering tasks; not relevant to this study) or used interviews or surveys to evaluate tools. We also removed papers that have a narrow focus or are entirely model-centric, e.g., interviewing only data scientists about their modeling work (e.g., [24, 36, 47, 81]) or interviewing only non-technical people (e.g., [13, 34, 101, 118]).

Snowballing. To capture relevant papers that did not match our keywords in their abstract, we performed one iteration of backward snowballing [127], which means that we went through the selected papers' reference list to find whether we missed any relevant papers. We analyzed 26 additional papers and considered 11 of them as relevant based on the inclusion and exclusion criteria, which included the three papers from the seed set we previously missed.

Final paper set. Overall, our process resulted in a final set of 50 papers. Most of the papers were published recently, since 2019 (see Figure 2). This sudden explosion of interview and survey studies with practitioners in recent years justifies our motivation for this study to aggregate all the findings of these papers. Most of the papers, 30 out of 50, were published in software engineering venues (including five at WAIN/CAIN), 11 papers in HCI venues, two papers in AI Ethics venues, and the seven remaining ones are scattered over other communities. A total of 947 interviews and 3811 survey

**Figure 2: Year Distribution of the Selected Papers**

responses were reported in 43 papers, and the seven remaining papers did not report specific counts of the interviewed or surveyed practitioners.

Of the 50 papers, 31 papers explicitly list research questions or the aim of their research as identification of challenges (or issues, problems, difficulties) in different aspects of building products with ML components. The other papers do not explicitly set a goal of identifying challenges but more broadly study the process of building products with ML components, yet they also report practitioner challenges in their findings.

3.2 Qualitative Meta-Summary Process

As stated earlier, we used the meta-summary research method [25, 96, 104] to synthesize the findings from the collected papers. This method is used to perform quantitative aggregation of qualitative findings, which are necessarily the thematic summaries of the underlying data from different studies. We conduct the following steps to perform the synthesis, as per the guidelines.

Extracting findings. Along with the standard metadata (title, source, venue, year, etc.), we extracted study-specific data regarding research questions, study method, interview and survey participant counts, and, most importantly, the challenges reported within the papers. To maintain consistency in extracting the findings, we considered only *challenges* that were derived from the interview and survey answers in the papers, not challenges derived from other literature or personal experience of the authors. We extracted challenges related to building software systems with ML components, but excluded those that relate exclusively to the data- and model-related work performed by a data scientist, such as algorithmic problems, notebook coding, and hyper-parameter tuning. We extracted a total of 520 excerpts relating to challenges from the 50 papers. We stored all extracted information from each paper in a spreadsheet for further analysis.

Grouping topically similar findings. We organize the findings at the level of *reported challenges* that we extracted from the papers. Different papers grouped findings in different ways and using different terminologies; we aimed to find a consistent organizational principle. For identifying similar topics and grouping those together, we needed to understand and compare those reported challenges in their original context. *Card sorting* is a common technique for grouping similar findings [41, 116], which we used for this paper. Following the standard *card sorting* method, we created one (virtual) card per reported challenge, and incrementally and iteratively organized those cards into groups of similar challenges. Multiple researchers went through all the cards in synchronous and asynchronous fashion to grasp the different concepts and identify

relevant themes and clusters around the reported challenges. This being a collaborative effort, we did not aim for inter-rater agreement between independent grouping by individual researchers, but instead worked together as a team to build consensus. There were many rounds of card sorting including moving the cards back and forth between different clusters, splitting the cards to handle different dimensions, merging similar clusters, and splitting clusters when we found there was more than one theme, until all involved researchers were satisfied with the clusters and placement of the cards. We developed three layers of clustering – the reported challenges extracted from the papers as the smallest unit, groups of common *themes* or patterns in the challenges as the second layer, and finally a third (or top) layer grouping the second layer clusters by development stages or cross-cutting concerns for the ease of reporting results. We performed this card sorting process in an online platform (miro.com), allowing us to manipulate colors, add different tags to the cards, add comments, emojis, and so on. We share the resulting card-sorting board as supplementary documents [77].

Abstracting and formatting findings. For each of the second layer clusters we abstracted out the concrete details of the reported challenges and summarized the clusters based on the identified themes of the groups. For this, we once again looked into the cards of each of the clusters individually and attempted to develop broad statements that capture the content of the cards in that cluster, which provide the headings of our results presented in Section 4. We wanted to be concise, but also comprehensive to properly capture the themes in the card. At the same time, as Sandelowski and Barroso suggested [15], we were careful to preserve the context in which the findings appeared by going back and forth in the original papers when confusion arose, moving cards to other clusters or themes as needed.

Calculating effect sizes. Methods for meta-summaries recommend reporting the frequency of findings in the original sources [104]. Since many of our analyzed papers ask similar broad research questions, we can carefully interpret findings mentioned more frequently as more common, though some papers clearly specialize in specific sub areas such as fairness or software architecture [37, 59]. We do not attempt to count frequencies of mentions within the papers (“intensity effect size”) because they are not consistently reported, but just report the percentage of papers reporting on a challenge theme (“frequency effect size”).

3.3 Limitations and Threats to Validity

All research designs come with limitations that threaten validity and credibility of results. As usual, readers should be careful when generalizing findings beyond what is allowed by the methods. Despite best efforts in our selection methods (SLR process, snowballing) we may have missed some relevant papers. In setting clear rules for scope, we had to do some judgment calls by consensus of all researchers for a number of papers, for example, whether to include [3, 7, 12, 36, 109].

As discussed earlier, the meta-summary synthesis method was chosen deliberately for its fit, but comes with its own limitations: it does not analyze original raw data, but only what is reported by other papers. Organizing and categorizing the data required

some interpretation of the papers and some judgment calls. The method encourages quantification of effect sizes, but those may not be entirely reliable as the analyzed papers use different methods and sometimes focus on specific subquestions.

It would have been interesting to analyze findings in additional dimensions, for example, whether team members in different roles or projects, or in different application domains, experience different challenges, or whether different challenges surface depending on the research method in the original study (e.g., survey vs. interview, open question vs. closed question). Unfortunately, data in the original studies is frequently not reported consistently and with enough granularity to enable such analyses.

While the meta-summary method can in principle also identify conflicts within the literature, this was not feasible in our study. The analyzed papers typically reported challenges, not the absence or relative importance of certain challenges. Given that different papers often had a different focus, rather than being replications of each other, we cannot conclude that not mentioning a challenge implies that there was no such challenge. Hence, we limited our analysis to aggregating and grouping reported challenges.

4 RESULTS

We report our findings of the meta-summary in this section using the layers derived from the card sorting. The top layer includes development stages (1) *Requirements Engineering*, (2) *Architecture, Design, and Implementation* (with a special focus on (2a) *Model Development* and (2b) *Data Engineering*), and (3) *Quality Assurance*, plus (4) *Process* challenges and (5) *Team* challenges as crosscutting concerns. Also, although *MLOps*, *Fairness*, and other more specific categories are often used to organize results in the surveyed papers, we eventually settled on minimizing the number of cross-cutting topics. We decided to include operations challenges in the *Architecture and Design* group, as we consider them primarily as a *design for change* issue; and we separate and group various concerns for specific qualities, such as *fairness*, in the development stages where the concerns arise, such as requirements and quality assurance. Within these top layer headings, we have our second layer clusters which are the abstracted challenges based on our identified themes, reported as the sub-headings in the following sections.

4.1 Requirements Engineering

Requirements engineering is known as an important and challenging stage of any software project, but as a consistent theme, we find that practitioners argue that the incorporation of ML further complicates requirements engineering.

Lack of AI literacy causes unrealistic expectations from customers, managers, and even other team members [7, 23, 38, 45, 52, 56, 62, 68, 78, 79, 86, 100, 110, 119, 120, 123, 128] (17/50).

Across many studies, many practitioners report that customers frequently have unrealistic expectations of ML capabilities in a product, like demanding a complete lack of false positives or expecting very high accuracy that is infeasible with provided resources (e.g., data, funding). Commonly, practitioners similarly blame a lack of *AI literacy* on customers not wanting to pay for the continuous improvement of the model: they have a static view of model development [45, 78] only consider paying for coding, as they do

not understand the need for experimental analysis [62] and even difficulty convincing engineering teams to invest in collecting high-quality data [52]. The issue of unrealistic requirements does not only come from customers, but also from team members within the company itself: Data scientists find it hard to explain the capabilities of ML to managers, requirements engineers, and even designers [23, 38, 78, 79, 86, 123]. According to practitioners, a lack of AI literacy in team members manifests particularly in defining and scoping the project: Stakeholders find it hard to understand the suitability of applying ML itself [56, 128], scoping and deciding the functional and non-functional requirements [62, 120], interpreting the model outcomes [79, 100, 110], and the infrastructure needs (e.g., appropriate data, monitoring infrastructure, retraining requirements) when building products [79, 120, 128]. Many practitioners also report that ML-specific system-level qualities like fairness and explainability are frequently ignored during requirements elicitation, as the stakeholders are not aware of them [11, 78, 95, 120].

Vagueness in ML problem specifications makes it difficult to map business goals to performance metrics [30, 37, 56, 58, 61, 62, 66, 78, 79, 86, 95, 100, 110, 115, 119, 120, 123] (17/50). Practitioners across many studies mention the challenge of formulating the specific software and ML problem in a way that satisfies business goals and objectives. ML practitioners find it difficult to map the high-level business goals to the low-level requirements for a model. While customers are broadly interested in improving the business, practitioners often find it difficult to quantify the contribution of the ML model and its return on investment. Also, *Responsible AI* initiatives find it difficult to quantify their contributions to the business, for example, measuring the value added by improving fairness and explainability, or to deliberate about tradeoffs between conflicting fairness and business objectives [11, 37, 86, 95]. Even with some notion of the responsible AI requirements in hand, practitioners find the requirements vague and not concrete enough to actually implement (e.g., unclear subpopulations and protected characteristics to balance discrimination) [95, 120]. On the other hand, practitioners also frequently report that many projects are exploratory without clear upfront business goals, thus, starting off the project without clear requirements is pretty common, albeit often problematic [30, 62, 110, 123].

Regulatory constraints specific to data and ML introduce additional requirements that restrict development [11, 30, 38, 82, 109, 110, 120] (7/50). Practitioners in multiple studies expressed how regulatory restrictions constrain ML development and require audits and involvement from legal teams. Privacy laws such as GDPR impose additional requirements on ML practitioners such as ensuring the collection of individual consent [38, 120] and providing the nontrivial ability to remove individuals from training data after they revoke consent. Similarly, practitioners in regulated domains report a need for explainability and transparency that prevents them from using deep learning and post-hoc explainability techniques [11, 30, 109].

4.2 Architecture, Design, and Implementation

We find that many ML practitioners struggle with designing the architecture of products with ML components.

Transitioning from a model-centric to a pipeline-driven or system-wide view is considered important for moving into production, but a difficult paradigm shift for many teams [1, 43, 56, 59, 62, 66, 68, 74, 76, 110, 128] (11/50). Practitioners frequently report challenges in migrating from exploratory model code, often in a notebook, to deployable production-quality code in automated ML pipelines [62, 128]. Building an end-to-end ML pipeline is considered to be a challenge due to the difficulties of integrating various ML and non-ML components in a system operating within an environment [56, 59, 110], the overwhelming complexity of integrating many tools and frameworks [66, 68, 74], the need for engineering skills beyond the comfort zone of some data scientists [74], and so on. While practitioners emphasized the importance of pipeline automation for many projects where frequent re-training and deployment of models are needed, they also consider it time-consuming, labor-intensive, error-prone, and not well supported by current tools [1, 43, 66, 76, 110].

ML adds substantial design complexity with many, often implicit, data and tooling dependencies, and entanglements due to a lack of modularity [1, 5, 23, 59, 61, 66, 110, 111, 115, 123, 128] (11/50). Many practitioners report challenges from additional complexity when designing systems incorporating machine learning, and the traditional software architecture and design practices no longer fit [23, 59, 61, 128]. ML changes the assumptions in traditional software systems such as encapsulation and modularity and causes entanglements of data, source code, and ML models, which can lead to “*pipeline jungles*” and “*change anything changes everything*” integrations that are hard to maintain [1, 61, 66, 110, 111, 123]. Unlike traditional systems, ML requires the incorporation of data pipelines that need to handle a high volume of data and often data architectures of distributed nature, and practitioners also need to understand and design for the data flow in the entire system [115, 123, 128]. Practitioners also point out that complexities arise due to a large amount of surrounding “glue code” to support the ML models [5, 110], and complicated dependency and configuration management [5, 123].

Difficulty in scaling model training and deployment on diverse hardware [30, 43, 62, 66, 68, 74, 100, 108, 110, 111] (10/50). Practitioners commonly report difficulty dealing with cloud and computational resources, even with the recent emergence of MLOps. Practitioners find the technologies to be difficult to integrate into the production environment and require substantial time, effort, and money [30, 62, 66, 68, 74, 100]. Among the common problems of such deployments, practitioners brought up the mismatch of development and production environments [62, 111], difficulties in building a scalable pipeline [30, 43, 66, 108], adhering to serving requirements such as latency and throughput [66, 110], as well as undocumented tribal knowledge within the team, hampering future deployments [111]. Despite the emerging MLOps tooling, practitioners still raise many questions about how to utilize those resources and sometimes express being overwhelmed by the sudden flood of tools and frameworks to choose from [3, 52].

While monitorability and planning for change are often considered important, they are mostly considered only late after launching [1, 5, 11, 30, 43, 55, 56, 58, 59, 74, 78, 99, 110, 111,

128] (15/50). Practitioners report struggling with monitoring their deployed models for detecting drift, bias, or even failures. While many highlight monitoring as very important, planning for monitoring is rare [78]. Even for companies that adopt a monitoring infrastructure, practitioners report struggling with ad-hoc monitoring practices of logging, creating alerts, or doing everything manually [59, 111]. Similar concerns were raised about model evolution, where practitioners acknowledge it to be important, but fall behind in planning for change in their architectural design [1, 5, 30, 43, 110]. Practitioners mentioned that ML-centric software goes through frequent revisions more than traditional software (e.g., due to model retraining, or even model replacement for data change, hyperparameter tuning, or change of domain, etc.), and the changes tend to be nontrivial and nonlocal, raising the need for an architecture that supports such changes. As a result, we find practitioners' soliciting the need for adapted architectural patterns to design for such post-launch activities for products with ML components with monitorability as a significant quality attribute [59, 128].

4.3 Model Development

Although we explicitly exclude challenges relating only to the work and tools of data scientists when building models, we find reports of engineering challenges during model development, which we report in this section.

Model development benefits from engineering infrastructure and tooling but provided infrastructure and technical support are limited in many teams [3, 5, 8, 12, 29, 30, 43, 52, 55, 56, 58, 62, 76, 79, 82, 93, 119, 123, 131] (19/50). ML practitioners share tooling needs for different tasks including data analysis and visualization, feature engineering, model development, integration, evaluation, deployment, monitoring, reproducibility, and support for specific qualities like privacy, security, and explainability. They report a lack of adequate tools in these areas and find the existing tools and techniques to be (a) unavailable in their environment [8, 30], (b) not automated enough [93], (c) requiring too much expert knowledge to be used [58, 82, 93, 131], (d) limited to specific tasks and types of data sets [8, 93], or (e) not suitable for their own problems [8, 12, 52]. This raises demand for custom tools but many teams lack the resources and engineering support.

Code quality is not standardized in model development tools, leading to conflicts about code quality [78, 111, 123] (3/50). Practitioners report that code quality and review processes are usually not standardized and are inconsistent across development and production environments. The expectations around code quality and versioning also differ widely in teams and create conflicts within teams, especially among team members with different roles and backgrounds. Practitioners commonly complain about low code quality in data science code, especially in notebooks.

4.4 Data Engineering

In developing machine learning models, data plays an important role. While we exclude challenges related exclusively to data-related work within ML pipelines, we report engineering challenges related to handling data within the system.

Data quality is considered important, but difficult for practitioners and not well supported by tools [1, 5, 29, 30, 38, 52, 61, 62, 66, 68, 74, 78, 93, 100, 110, 111, 120] (17/50). ML practitioners commonly report struggling with validating and improving data quality. Even with significant research efforts in building tools for data labeling, cleaning, visualization, and management, data work is still reported as a problematic area for practitioners. Practitioners reported that they need to invest significant effort and time in data pre-processing, cleaning, and assembly [2, 29, 38, 52, 62, 66, 78, 93]. Practitioners also mention their pain points in handling data errors and validating data quality, where better tool support is desired [52, 61, 74, 100, 108, 110, 111, 120]. Although it is common to associate these data issues within the model building pipeline, practitioners feel the need for cooperation from other parts of the organization (e.g., requirements engineers need to identify and specify requirements regarding data collection, formats, and the ranges of data and domain experts need to help to understand the structure and semantics of the data), which they mention is lacking [38, 78, 93, 100, 120].

Internal data security and privacy policies restrict data access and use [5, 30, 48, 52, 56, 62, 66, 68, 78, 100] (10/50). Data access is often restricted due to security and privacy policies within organizations, beyond possible regulatory restrictions, e.g., policies ensuring that customer data is not shared outside the company. Due to restrictions on the flow of data, ML practitioners need to deal with additional complexities in the data pipeline, as only a restricted number of team members can analyze the data and as they have limited access to the right data and no access to data locally for model optimization or model debugging due to data movement constraints [30, 48, 62, 66].

Although training-serving skew is common, many teams lack support for its required detection and monitoring [5, 30, 58, 66, 78, 111, 128] (7/50). The mismatch between training data and production data is a common problem in products with ML components, where models work well on test data but generalize poorly to real-world data in production. Even if training the model with a representative dataset initially, the production environment often encounters drift toward data distributions that are less well supported by the model. Practitioners explain that monitoring models in production for staleness is an important activity that supports detecting the degradation of model performance and retraining it with new data if needed. However, they also find it challenging to set up the monitoring infrastructure and report a lack of tool support.

Data versioning and provenance tracking are often seen as elusive, with not enough tool support [1, 38, 43, 56, 68, 108, 119] (7/50). While software engineers routinely adopt mature version control systems for code, practitioners report challenges in versioning data, typically due to the large volumes of data involved. Practitioners mention that they need to have traceability and transparency to answer questions like "Which data was this model trained on?" or "Which code or data change made our accuracy deteriorate?" [43], but it's not possible for them to keep track of data and models across the life cycle without technological support [2, 43, 129]. This is a bigger problem for practitioners in small companies as they do

not want to invest in storage capacity to version their models and datasets, though they understand the importance [38].

4.5 Quality Assurance

One of the biggest changes that the incorporation of ML models has brought into traditional software development is challenging the traditional notion of correctness, where models are evaluated for accuracy or fit rather than whether they fully meet a specification. Understandably this impacts the conventional processes and practices associated with testing and quality assurance.

Testing and debugging ML models is difficult due to lack of specifications [2, 5, 29, 30, 45, 58, 61, 62, 66, 76, 78, 93, 100, 108, 110, 111, 119, 123, 131] (19/50). Practitioners find testing and debugging of ML models challenging. In particular, they ubiquitously report difficulty establishing quality assurance criteria and metrics, given that no model is expected to be always correct, but it is difficult to define what amount and what kind of mistakes are acceptable for a model [5, 29, 30, 45, 61, 62, 66, 78, 123, 131]. In particular, practitioners find it difficult to define accuracy thresholds for evaluations. Furthermore, practitioners report finding it difficult to select adequate test data, specifically curating test data of sufficient quality and quantity that is representative of the production environment [58, 76, 93, 119, 123]. Curating test data for ML testing is also considered costly and labor-intensive, and practitioners desire methods and tools from the research community for automated test input generation to reduce this cost [45, 61, 123]. Practitioners consider it a challenge to get labels for test data and evaluate test quality (e.g., in terms of coverage) due to the difficulty of defining the valid input space and the test oracle problem [29, 61, 123]. Practitioners also mention the silent failing of models (i.e., models give wrong answers rather than crashing), the long tail of corner cases, and the “invisible errors”, that are handled on an ad-hoc basis without a systematic framework or a standard approach [29, 111, 123]. Additionally, practitioners raise challenges regarding evaluating model robustness, on one hand, suffering from the lack of a concrete methodology [29, 93], and on the other hand, having various metrics but no consensus on which metric to use [61].

Testing of model interactions, pipelines, and the entire system is considered challenging and often neglected [29, 56, 61, 66, 76, 78, 93, 131] (8/50). Testing literature often focuses on ML models and data quality, but less on how models are integrated into the system, and even less on the infrastructure to produce the models. Practitioners find sole unit testing of individual models insufficient and ineffective, due to the entanglement of models and different ML components, as well as the difficulty of explaining why an error occurred due to the low interpretability of individual models [61, 123, 131]. The lack of pipeline and system testing beyond the model is also considered a problematic area [29, 56, 76, 78, 93, 131]: While practitioners tend to focus more on the data- and model-related issues, the error handling around the model is found to be insufficient in previous studies [29, 131], leading to system failures even where the model gives the correct results [76]. Practitioners also report having no systematic evaluation strategy nor automated tools and techniques for pipeline and system-level testing [61, 78].

Testing and monitoring models in production are considered important but difficult, and often not done [61, 78, 93, 110, 111] (5/50). Many practitioners recognize the need to test in production (online testing), since offline test data for models may not be representative, especially as data distributions drift. However, practitioners consider online testing complex as it is not trivial for them to design online metrics that do not only depend on the model but also on the external environment, user interactions after deployment, and the context of the product overall [61, 111]. Practitioners also find online testing very time-consuming, as it requires longer observation periods to determine meaningful results [110, 111]. Practitioners also pointed out that there is no surefire strategy to precisely detect when the model is underperforming in online testing [111].

There are no standard processes or guidelines on how to assess system qualities such as fairness, security, and safety in practice [11, 12, 37, 38, 43, 55, 99, 109, 119] (9/50). Research often discusses how machine learning influences fairness, robustness, security, safety, and other qualities, but practitioners report that they find evaluating these as challenging. While practitioners consider these qualities important [43, 55], they often report having no effective methodology or concrete guidelines for evaluating them [12, 37, 38, 55, 99, 109, 119]. Even regarding fairness, which has received a lot of research attention lately, practitioners report finding it hard to apply auditing and de-biasing methods due to not having a proper process in place [37, 38]. Some practitioners report waiting for complaints from customers rather than being proactive when it comes to fairness [37], or even blindly expecting the algorithms to inherently provide qualities like security against attacks [55].

4.6 Process

Building software products with ML components involves many moving parts that need to be planned and integrated. Fitting all of these together in a cohesive process can be challenging.

Development of products with ML component(s) is often ad-hoc, lacking well-defined processes [5, 7, 30, 45, 52, 56, 62, 76, 115, 119, 123] (11/50). Many practitioners report that they struggle with finding a good process for developing ML components and products around them [30, 45, 62, 115, 123], often coming up with ad-hoc strategies and experiencing a lack of good engineering practices [45, 76]. ML practitioners have explored using the traditional software development life cycles and found those to be a poor fit for exploratory development work. Even with a flexible agile methodology, practitioners identified that small iterations of sprints cannot fit the initial feasibility study that ML requires, with the timeline being too fixed and too short [7, 30, 62]. Also, they find it hard to set expectations for each sprint, as the project objectives may remain unclear at the beginning and need to be revisited after the initial investigation [7, 56].

The uncertainty in ML development makes it hard to plan and estimate effort and time [5, 7, 29, 45, 62, 119, 123] (7/50). Machine learning work tends to be iterative and exploratory and as such uncertain, where practitioners cannot estimate upfront how long it may take to reach a model with a certain level of accuracy

or whether that is even possible at all; instead, they commonly progress with many experiments with different algorithms and datasets [5, 45]. Practitioners, therefore, report having difficulties setting expectations and (intermediate) deadlines for a project [5, 7, 29, 62, 123] and providing any upfront estimates about effort and cost [5, 62, 119].

Practitioners find documentation more important than ever in ML, but find it more challenging than traditional software documentation [19, 30, 54, 58, 62, 78, 89, 108, 129] (9/50). Many practitioners point out various process and coordination challenges rooted in poor documentation. Some practitioners emphasize that documentation is even more important when it comes to ML components, as human decisions are inscribed in different stages of ML pipelines and cannot be retrieved from code or data without documentation [30, 129]. The final model code is the outcome of many different explorations and experimentations that include multiple rounds of data processing, feature engineering, hyperparameter tuning, and other activities. Many problem-specific decisions have been made in those stages that cannot be understood from the resulting model or pipeline code. Some argue that not recording these decisions in documentation causes them to slowly become invisible, severely impacting future re-analysis and revisions, or even model integration and deployment [30, 58, 129]. Others emphasize that, along with model documentation, data documentation is also imperative to share hidden information inside the data and create a shared data understanding, yet mostly missing in organizations [78, 108]. Others report that, with the incorporation of ML, the documentation process becomes more complicated as ML practitioners find it difficult to present complex model information in an accessible way to all levels of stakeholders [19, 54, 89]. It is also non-trivial for practitioners to decide on the right amount of details to include in the documentation. They place the blame mostly on the lack of organizational incentives, resources, and unclear and vague guidelines for ML documentation [19, 89].

4.7 Organization and Teams

Along with the challenges faced in different development stages, practitioners also mention challenges they suffer from the organizational and teamwork perspective while building products with ML components.

Building products with ML components requires diverse skill sets, which is often missing in development teams [3, 5, 7, 48, 68, 76, 79, 110, 119, 123, 124, 128] (12/50). Incorporation of ML in a product does not merely mean adding just another component to the system; it requires people from multiple disciplines to get involved to support different aspects of this component. The team requires many diverse skill sets to develop, deploy, and integrate the model into the complete product, including hardware expertise, engineering skills, knowledge of math and statistics, business understanding, UX design ability, operations, and domain expertise. The lack of this varied expertise in the team is commonly mentioned to be a challenge by practitioners [3, 7, 48, 110, 119, 123]. Also, as discussed in the next subsection that communication is often hindered by a lack of AI literacy or common terminology [1, 5, 30, 76, 79, 128], cross-disciplinary knowledge seems to be important for team members to interact and understand each other's

vocabulary; however, practitioner experiences seem to indicate that such cross-disciplinary education is not broadly available yet [29, 111].

Many teams are not well prepared for the extensive interdisciplinary collaboration and communication needed in ML products [5, 8, 18, 23, 68, 76, 78, 79, 108, 123, 128] (11/50). For building a product with ML components, team members need to collaborate with people from different disciplines as mentioned above, such as business leaders, engineers, designers, and various other departments inside the company, and even outside the organization [23, 79, 108, 123, 128]. Practitioners report that they often struggle to collaborate effectively in such interdisciplinary teams, because team members often do not understand the concerns of other members from other backgrounds, like data scientists lacking knowledge of engineering practices, testing frameworks, continuous integration and delivery, and such [1, 5, 30, 76]; software engineers lacking AI literacy [30, 76]; and data scientists and software engineers not understanding or interacting members with in with business roles [79, 128]. Practitioners report struggling with cultural differences, differences in expectations, and conflicting priorities [5, 8, 68, 79], and they often do not agree on assigned responsibilities [62, 78]. These multidisciplinary teams also suffer from miscommunications arising from inconsistency in their technical terminologies [18, 76, 78]. Siloing of teams by specialization and lack of communication across such silos are also observed in many production settings, fostering integration problems even further [8, 78].

ML development can be costly and resource limits can substantially curb/limit efforts [3, 5, 48, 68, 79, 111] (6/50). Practitioners report that organizations involved in the development of products with ML components often suffer from resource and budget limitations. Hardware, infrastructure, cloud storage, GPUs, etc., are expensive, and especially for small companies, it is difficult to justify such expenditures based on the expected return on investment from the model.

Lack of organizational incentives, resources, and education hampers achieving all system-level qualities [12, 48, 55, 79, 82, 95, 99, 109] (8/50). Practitioners mention that organizational incentives also have an impact on achieving certain qualities of products with ML components. A quality that practitioners reported frequently as particularly challenging due to the lack of organizational incentives is fairness [37, 95]. Awareness of potential problems, including potential consequences from biased models, seems to be the main reason for lacking responsible AI practices, along with the lack of organizational incentives and structures, as well as priority conflicts. Safety, security, and privacy also seem to suffer from similar issues of awareness, education, resource constraints, and are often disregarded due to tradeoffs with development cost [12, 48, 55, 82, 99].

5 DISCUSSION AND CONCLUSIONS

With this meta-summary, we aggregate and summarize the challenges reported by industry practitioners who build software products with ML components. We find that practitioners report challenges in all stages of the development process, from the initial

requirements specification stage to quality assurance of the deployed product. They report a broad range of issues from lacking process, organizational structure, and team collaboration strategies, to lacking tool support for data, model building, deployment, and monitoring.

Old, new, and harder challenges. Arguably, many reported challenges are not new to software engineers, and likely many software engineers may have reported similar challenges in non-ML projects. It seems though that the introduction of machine learning exacerbates some universal challenges and introduces new ones. For example, software engineering literature is well aware that requirements engineering is challenging, with customers having unrealistic expectations and developers directly jumping into coding without understanding requirements first. While our study does not support direct comparisons, it seems that these problems haunt ML practitioners more, given how ML inspires hopes for amazing capabilities, but in a way that may be difficult to understand and specify without substantial ML expertise. Similarly, the software-engineering literature is full of nuanced discussions of development life cycles and competing process models, but ML practitioners struggle adopting even the most flexible agile-inspired processes for their projects with the uncertainty that ML brings. Also, team collaboration and organizational challenges are well known in traditional software engineering, but those seem to become even more central with the additional complexity and inclusion of more people with different backgrounds, cultures, and priorities. Other challenges seem new, such as the data- and model-related challenges associated with ML components, and several of the reported challenges regarding architecture and quality assurance stemming from the different nature of reasoning in machine learning.

Toward better engineering of ML products. A finding from our study is that there is much more consensus on what the challenges are, than how to overcome them. Some challenges could be addressed with new tooling or new practices; for others it may be possible to simply adopt existing good engineering practices; and yet others may just be intrinsically hard problems. While we cannot provide a rigorous summary or analysis, we close by reflecting on possible directions.

- **Requirements Engineering.** For the challenges of unrealistic requirements, several studies mentioned that practitioners found it useful to conduct training sessions with clients and other team members on AI literacy, before starting the ML projects [78, 96, 107, 120]. But again, while many practitioners mention suffering from unclear model requirements, we still do not seem to have a good solution to that, and additional research on how to elicit and describe requirements for models may be needed. Another area for future research would be to better understand and prepare for regulatory constraints and provide evidence of compliance.
- **Architecture, Design, and Implementation.** Machine learning seems to provide significant challenges to architectural design of software systems, but arguably many challenges are similar to other large and complex and distributed software systems. While there are nascent discussions on organizing architecture knowledge as patterns [57, 59, 110, 125], it does not seem like the field has reached saturation. This

seems to be a field though, where industry-oriented research (similar to the data architecture of facebook [32]) has more access to the complicated real-world scenarios where architectural planning becomes important than what academics can typically access. From the challenges raised by practitioners, it is apparent that along with the need for design practices, patterns, and mechanisms to handle system and model-level considerations (e.g., dependency management, scalability, monitorability), we also need to support teams in shifting from model-centric work to system thinking, possibly through tailored education for ML practitioners.

- **Model Development and Data Engineering.** Consistent across many papers, we find that ML practitioners desire more engineering support, such as better infrastructure and tools for model and data work. Data scientists also indicate a need for more cooperation from other team members in terms of support for data, which necessitates better collaboration strategies and data education for the entire team. On the other hand, a few practitioners highlighted the necessity of standardization of ML code quality, which may be a low hanging fruit technically, but may require a change to the culture and practices in many projects.
- **Quality Assurance.** Quality assurance for machine learning, especially for models, is a very active area of research, with proposals for many different testing strategies to validate different model characteristics covered in multiple literature surveys [39, 97, 98, 130]. While we found that a lot of practitioners mentioned concerns about specifying model adequacy goals, few practitioners showed concerns about system testing, monitoring in production, and testing for fairness, security, and safety. We are surprised to not see more concerns about system-level quality beyond the model, which might indicate either that practitioners do not consider these testing areas as challenging, or that most organizations (especially outside of big tech) are not yet mature enough to even start thinking about such testing needs. Monitoring though is recognized as an important challenge, with many available tools but common adoption problems that may be worth investigating further.
- **Process.** While there is research on the development processes for ML models [70, 117], there seems to be little work on addressing process challenges that arise when integrating ML and non-ML work in production projects that are commonly mentioned by practitioners. We believe that this is an area with plenty of research opportunities to evaluate what processes and practices work well in different contexts.
- **Organization and Teams.** While there is lots of research on technical issues, practitioners often see organizational and team issues (such as a lack of AI literacy in teams, unclear responsibility boundaries, and a lack of team synchronization) as some of the most difficult challenges to overcome. Education and better collaboration strategies seem to be the factors that might put a positive impact on mitigating many of the challenges that the practitioners mentioned.

Overall we believe that a lot of progress can be made with better education and better adoption of good software engineering practices. There are plenty research opportunities to adapt existing

practices, support them with tooling, and create new interventions altogether. We hope that the collection of challenges, which can be traced to the original studies where they were raised by practitioners, will be helpful in selecting and prioritizing research and education in our community.

Acknowledgments. Kästner’s, Nahar’s, and Zhang’s work was supported in part by the National Science Foundation (#2131477), Zhou’s work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC, RGPIN2021-03538), and Lewis’ work was funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center (DM23-0228).

REFERENCES

- [1] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B. and Zimmermann, T. 2019. Software Engineering for Machine Learning: A Case Study. *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)* (2019), 291–300.
- [2] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B. and Zimmermann, T. 2019. Software Engineering for Machine Learning: A Case Study. *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)* (2019), 291–300.
- [3] Andrade, H., Lwakatere, L.E., Crnkovic, I. and Bosch, J. 2019. Software Challenges in Heterogeneous Computing: A Multiple Case Study in Industry. *Proceedings of the 45th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (2019), 148–155.
- [4] Arnold, M., Piorowski, D., Reimer, D., Richards, J., Tsay, J., Varshney, K.R., Belamy, R.K.E., Hind, M., Houde, S., Mehta, S., Mojsilovic, A., Nair, R., Ramamurthy, K.N. and Olteanu, A. 2019. FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM journal of research and development*. 63, 4/5 (2019), 6:1–6:13.
- [5] Arpteg, A., Brinne, B., Crnkovic-Friis, L. and Bosch, J. 2018. Software Engineering Challenges of Deep Learning. *Proceedings of the 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (2018), 50–59.
- [6] Ashmore, R., Calinescu, R. and Paterson, C. 2022. Assuring the Machine Learning Lifecycle: Desiderata, Methods, and Challenges. *ACM Computing Surveys*. 54, 5 (2022), 1–39.
- [7] Baijens, J., Helms, R. and Iren, D. 2020. Applying Scrum in Data Science Projects. *Proceedings of the 22nd Conference on Business Informatics (CBI)* (2020), 30–38.
- [8] Bäuerle, A., Cabrera, Á.A., Hohman, F., Maher, M., Koski, D., Suau, X., Barik, T. and Moritz, D. 2022. Symphony: Composing Interactive Interfaces for Machine Learning. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022), 1–14.
- [9] Begel, A. and Zimmermann, T. 2014. Analyze this! 145 questions for data scientists in software engineering. *Proceedings of the 36th International Conference on Software Engineering* (2014), 12–23.
- [10] Bernardi, L., Mavridis, T. and Estevez, P. 2019. 150 Successful Machine Learning Models: 6 Lessons Learned at Booking.com. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining - KDD '19* (2019), 1743–1751.
- [11] Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J.M.F. and Eckersley, P. 2020. Explainable machine learning in deployment. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020), 648–657.
- [12] Boenisch, F., Battis, V., Buchmann, N. and Poikela, M. 2021. “I Never Thought About Securing My Machine Learning Systems”: A Study of Security and Privacy Awareness of Machine Learning Practitioners. *Proceedings of Mensch und Computer 2021* (2021), 520–546.
- [13] Borch, C. 2022. Machine learning, knowledge risk, and principal-agent problems in automated trading. *Technology in society*. 68, (2022), 101852.
- [14] Borg, M., Englund, C., Wnuk, K., Duran, B., Levandowski, C., Gao, S., Tan, Y., Kaijser, H., Lönn, H. and Törnqvist, J. 2020. Safely Entering the Deep: A Review of Verification and Validation for Machine Learning and a Challenge Elicitation in the Automotive Industry. *Journal of Automotive Software Engineering*. 1, 1 (2020), 1–19.
- [15] Boyd, K.L. 2021. Datasheets for Datasets help ML Engineers Notice and Understand Ethical Issues in Training Data. *Proceedings of the ACM on Human-Computer Interaction*. 5, CSCW2 (2021), 1–27.
- [16] Braiek, H.B. and Khomh, F. 2020. On testing machine learning programs. *The Journal of systems and software*. 164, (2020), 110542.
- [17] Breck, E., Cai, S., Nielsen, E., Salib, M. and Sculley, D. 2017. The ML test score: A rubric for ML production readiness and technical debt reduction. *Proceedings of the 2017 IEEE International Conference on Big Data (Big Data)* (2017), 1123–1132.
- [18] Brennen, A. 2020. What Do People Really Want When They Say They Want “Explainable AI?” We Asked 60 Stakeholders. *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), 1–7.
- [19] Chang, J. and Custis, C. 2022. Understanding Implementation Challenges in Machine Learning Documentation. *Equity and Access in Algorithms, Mechanisms, and Optimization* (2022), 1–8.
- [20] Chattopadhyay, S., Prasad, I., Henley, A.Z., Sarma, A. and Barik, T. 2020. What’s wrong with computational notebooks? Pain points, needs, and design opportunities. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), 1–12.
- [21] Chotisarn, N., Merino, L., Zheng, X., Lonpalawong, S., Zhang, T., Xu, M. and Chen, W. 2020. A systematic literature review of modern software visualization. *Journal of visualization / the Visualization Society of Japan*. 23, 4 (2020), 539–558.
- [22] Dilhara, M., Ketkar, A. and Dig, D. 2021. Understanding Software-2.0: A Study of Machine Learning Library Usage and Evolution. *ACM Transactions on Software Engineering and Methodology*. 30, 4 (2021), 1–42.
- [23] Dove, G., Halskov, K., Forlizzi, J. and Zimmerman, J. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 278–288.
- [24] Epperson, W., Wang, A.Y., DeLine, R. and Drucker, S.M. 2022. Strategies for Reuse and Sharing among Data Scientists in Software Teams. *Proceedings of the 44th International Conference on Software Engineering: Software Engineering in Practice* (2022), 243–252.
- [25] Faan, M.S.P. and Aprn, J.B.P. 2006. *Handbook for Synthesizing Qualitative Research*. Springer Publishing Company.
- [26] Felizardo, K.R., Mendes, E., Kalinowski, M., Souza, Ê.F. and Vijaykumar, N.L. 2016. Using Forward Snowballing to update Systematic Reviews in Software Engineering. *Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement* (2016), 1–6.
- [27] Follow, R. Bridging the Gap Between Data Science & Engineer: Building High-Performance Teams.
- [28] Giray, G. 2021. A Software Engineering Perspective on Engineering Machine Learning Systems: State of the Art and Challenges. *Journal of Systems and Software*. 180, (2021), 111031.
- [29] Golendukhina, V., Lenarduzzi, V. and Felderer, M. 2022. What is software quality for AI engineers? Towards a thinning of the fog. *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI* (2022), 1–9.
- [30] Haakman, M., Cruz, L., Huijgens, H. and van Deursen, A. 2021. AI Lifecycle Models Need To Be Revised. An Exploratory Study in Fintech. *Empirical Software Engineering*. 26, 5 (2021), 1–29.
- [31] Harris, J. 2020. Beyond the jupyter notebook: how to build data science products. *Towards Data Science*.
- [32] Hazelwood, K. et al. 2018. Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective. *Proceedings of the 2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)* (Feb. 2018), 620–629.
- [33] Head, A., Hohman, F., Barik, T., Drucker, S.M. and DeLine, R. 2019. Managing messes in computational notebooks. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19* (2019), 1–12.
- [34] Henry, K.E., Kornfield, R., Sridharan, A., Linton, R.C., Groh, C., Wang, T., Wu, A., Mutlu, B. and Saria, S. 2022. Human-machine teaming is key to AI adoption: clinicians’ experiences with a deployed machine learning system. *NPJ digital medicine*. 5, 1 (2022), 1–6.
- [35] Hermann, J. and Del Balso, M. 2017. Meet Michelangelo: Uber’s machine learning platform.
- [36] Hill, C., Bellamy, R., Erickson, T. and Burnett, M. 2016. Trials and tribulations of developers of intelligent systems: A field study. *Proceedings of the IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (2016), 162–170.
- [37] Holstein, K., Wortman Vaughan, J., Daumé, H., Dudik, M. and Wallach, H. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), 1–16.
- [38] Hopkins, A. and Booth, S. 2021. Machine Learning Practices Outside Big Tech: How Resource Constraints Challenge Responsible Development. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (2021), 134–145.
- [39] Huang, X., Kroening, D., Ruan, W., Sharp, J., Sun, Y., Thamo, E., Wu, M. and Yi, X. 2020. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Computer Science Review*. 37, (2020), 100270.
- [40] Huang, X., Zhang, H., Zhou, X., Babar, M.A. and Yang, S. 2018. Synthesizing qualitative research in software engineering: a critical review. *Proceedings of the 40th International Conference on Software Engineering* (2018), 1207–1218.
- [41] Hudson, W. 2013. Card Sorting. *The Encyclopedia of Human-Computer Interaction, 2nd Ed.* The Interaction Design Foundation.

- [42] Hulten, G. 2018. *Building Intelligent Systems: A Guide to Machine Learning Engineering*. Apress.
- [43] Hummer, W., Muthusamy, V., Rausch, T., Dube, P., El Maghraoui, K., Murthi, A. and Oum, P. 2019. ModelOps: Cloud-Based Lifecycle Management for Reliable and Trusted AI. *Proceedings of the 2019 IEEE International Conference on Cloud Engineering (IC2E)* (2019), 113–120.
- [44] Hynes, N., Sculley, D. and Terry, M. 2017. The data linter: Lightweight, automated sanity checking for ml data sets. *NIPS ML Sys Workshop*. 1, (2017), 5.
- [45] Ishikawa, F. and Yoshioka, N. 2019. How do engineers perceive difficulties in engineering of machine-learning systems? - questionnaire survey. *Proceedings of the 2019 IEEE/ACM Joint 7th International Workshop on Conducting Empirical Studies in Industry (CESI) and 6th International Workshop on Software Engineering Research and Industrial Practice (SER&IP)* (2019), 2–9.
- [46] Jain, R. and Suman, U. 2015. A Systematic Literature Review on Global Software Development Life Cycle. *SIGSOFT Softw. Eng. Notes*. 40, 2 (2015), 1–14.
- [47] Jentzsch, S. and Hochgeschwender, N. 2021. A qualitative study of Machine Learning practices and engineering challenges in Earth Observation. *it - Information Technology*. 63, 4 (2021), 235–247.
- [48] John, M.M., Olsson, H.H. and Bosch, J. 2020. AI Deployment Architecture: Multi-Case Study for Key Factor Identification. *Proceedings of the 27th Asia-Pacific Software Engineering Conference (APSEC)* (2020), 395–404.
- [49] Kanagal, B. and Tata, S. 2018. Recommendations for All: Solving Thousands of Recommendation Problems Daily. *Proceedings of the 34th International Conference on Data Engineering (ICDE)* (2018), 1404–1413.
- [50] Kästner, C. 2022. *Machine Learning in Production: From Models to Products*.
- [51] Keele, S. 2007. *Guidelines for performing systematic literature reviews in software engineering*. Technical Rep., Ver. 2.3 EBSE Tech. Report. EBSE.
- [52] Kim, M., Zimmermann, T., DeLine, R. and Begel, A. 2018. Data Scientists in Software Teams: State of the Art and Challenges. *IEEE Transactions on Software Engineering*. 44, 11 (2018), 1024–1038.
- [53] Kim, M., Zimmermann, T., DeLine, R. and Begel, A. 2016. The emerging role of data scientists on software development teams. *Proceedings of the 38th International Conference on Software Engineering* (2016), 96–107.
- [54] Königstorfer, F. and Thalmann, S. 2022. AI Documentation: A path to accountability. *Journal of Responsible Technology*. 11, (2022), 100043.
- [55] Kumar, R.S.S., Nystrom, M., Lambert, J., Marshall, A., Goertzel, M., Comisneru, A., Swann, M. and Xia, S. 2020. Adversarial Machine Learning - Industry Perspectives. *Proceedings of the 2020 IEEE Security and Privacy Workshops (SPW)*. (2020), 69–75.
- [56] Laato, S., Birkstedt, T., Määntymäki, M., Minkinen, M. and Mikkonen, T. 2022. AI governance in the system development life cycle: insights on responsible machine learning engineering. *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI* (2022), 113–123.
- [57] Lakshmanan, V., Robinson, S. and Munn, M. 2020. *Machine Learning Design Patterns*. O'Reilly Media, Inc.
- [58] Lewis, G.A., Bellomo, S. and Ozkaya, I. 2021. Characterizing and Detecting Mismatch in Machine-Learning-Enabled Systems. *Proceedings of the IEEE/ACM 1st Workshop on AI Engineering-Software Engineering for AI (WAIN)* (2021), 133–140.
- [59] Lewis, G.A., Ozkaya, I. and Xu, X. 2021. Software Architecture Challenges for ML Systems. *Proceedings of the 2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)* (2021), 634–638.
- [60] Lin, J. and Kolcz, A. 2012. Large-scale machine learning at twitter. *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data* (2012), 793–804.
- [61] Li, S., Guo, J., Lou, J.-G., Fan, M., Liu, T. and Zhang, D. 2022. Testing machine learning systems in industry: an empirical study. *Proceedings of the 44th International Conference on Software Engineering: Software Engineering in Practice* (2022), 263–272.
- [62] Liu, H., Eksmo, S., Risberg, J. and Hebig, R. 2020. Emerging and Changing Tasks in the Development Process for Machine Learning Systems. *Proceedings of the International Conference on Software and System Processes* (2020), 125–134.
- [63] Liu, J., Boukhelifa, N. and Eagan, J.R. 2020. Understanding the Role of Alternatives in Data Analysis Practices. *IEEE transactions on visualization and computer graphics*. 26, 1 (2020), 66–76.
- [64] Liu, Q., Li, P., Zhao, W., Cai, W., Yu, S. and Leung, V.C.M. 2018. A Survey on Security Threats and Defensive Techniques of Machine Learning: A Data Driven View. *IEEE Access*. 6, (2018), 12103–12117.
- [65] Lopez, G. and Guerrero, L.A. 2017. Awareness Supporting Technologies used in Collaborative Systems: A Systematic Literature Review. *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (2017), 808–820.
- [66] Lwakatare, L.E., Raj, A., Bosch, J., Olsson, H.H. and Crnkovic, I. 2019. A taxonomy of software engineering challenges for machine learning systems: An empirical investigation. *Proceedings of the 2019 International Conference on Agile Software Development* (2019), 227–243.
- [67] Lwakatare, L.E., Raj, A., Crnkovic, I., Bosch, J. and Olsson, H.H. 2020. Large-scale machine learning systems in real-world industrial settings: A review of challenges and solutions. *Information and software technology*. 127, 106368 (2020), 106368.
- [68] Mäkinen, S., Skogström, H., Laaksonen, E. and Mikkonen, T. 2021. Who Needs MLOps: What Data Scientists Seek to Accomplish and How Can MLOps Help? *Proceedings of the IEEE/ACM 1st Workshop on AI Engineering - Software Engineering for AI (WAIN)* (2021), 109–112.
- [69] Martínez-Fernández, S., Bogner, J., Franch, X., Oriol, M., Siebert, J., Trendowicz, A., Vollmer, A.M. and Wagner, S. 2022. Software Engineering for AI-Based Systems: A Survey. *ACM Transactions on Software Engineering and Methodology*. 31, 2 (2022), 1–59.
- [70] Martínez-Plumed, F., Contreras-Ochando, L., Ferri, C., Hernandez Orallo, J., Kull, M., Lachiche, N., Ramirez Quintana, M.J. and Flach, P.A. 2020. CRISP-DM twenty years later: From data mining processes to data science trajectories. *IEEE transactions on knowledge and data engineering*. 33, 8 (2020), 3048–3061.
- [71] McGlohon, M. 2021. Demystifying Machine Learning in Production: Reasoning about a Large-Scale ML Platform.
- [72] McGraw, G., Figueroa, H., Shepardson, V. and Bonett, R. 2020. An architectural risk analysis of machine learning systems: Toward more secure machine learning. *Berryville Institute of Machine Learning, Clarke County, VA*. Accessed on: Mar. 23, (2020).
- [73] Mitchell, M., Wu, S., Zaldívar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D. and Gebru, T. 2019. Model Cards for Model Reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2019), 220–229.
- [74] Muiruri, D., Lwakatare, L.E., K Nurminen, J. and Mikkonen, T. 2022. Practices and Infrastructures for ML Systems—An Interview Study in Finnish Organizations. *TechRxiv*.
- [75] Muller, M., Lange, I., Wang, D., Piorowski, D., Tsay, J., Liao, Q.V., Dugan, C. and Erickson, T. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), 1–15.
- [76] Myllyaho, L., Raatikainen, M., Männistö, T., Nurminen, J.K. and Mikkonen, T. 2022. On misbehaviour and fault tolerance in machine learning systems. *Journal of Systems and Software*. 183, (2022), 111096.
- [77] Nahar, N. 2022. Supplementary documents: A meta-summary of challenges in building products with ML components – collecting experiences from 4758+ practitioners. OSF.
- [78] Nahar, N., Zhou, S., Lewis, G. and Kästner, C. 2022. Collaboration Challenges in Building ML-Enabled Systems: Communication, Documentation, Engineering, and Process. *Proceedings of the 44th International Conference on Software Engineering* (2022), 413–425.
- [79] Namvar, M., Intezari, A., Akhlaghpour, S. and Brienza, J.P. 2022. Beyond effective use: Integrating wise reasoning in machine learning development. *International journal of information management*. (2022), 102566.
- [80] Nascimento, E., Nguyen-Duc, A., Sundbø, I. and Conte, T. 2020. Software engineering for artificial intelligence and machine learning software: A systematic literature review. *arXiv [cs.SE]*.
- [81] Nikanjam, A., Morovati, M.M., Khomf, F. and Ben Braiek, H. 2022. Faults in deep reinforcement learning programs: a taxonomy and a detection approach. *Automated software engineering*. 29, 1 (2022).
- [82] Nikhil, K., Anandayuvraj, D., Detti, A., Lee Bland, F., Rahaman, S. and Davis, J.C. 2022. "If security is required": Engineering and Security Practices for Machine Learning-based IoT Devices. *Proceedings of the 4th International Workshop on Software Engineering Research and Practices for the IoT (SER4IoT)* (2022), 1–8.
- [83] Nushi, B., Kamar, E., Horvitz, E. and Kossman, D. 2017. On human intellect and machine failures: troubleshooting integrative machine learning systems. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (2017), 1017–1025.
- [84] Ozkaya, I. 2020. What is really different in engineering AI-enabled systems? *IEEE software*. 37, 4 (2020), 3–6.
- [85] Paley, A., Urma, R.-G. and Lawrence, N.D. 2022. Challenges in deploying machine learning: A survey of case studies. *ACM computing surveys*. (2022).
- [86] Passi, S. and Jackson, S.J. 2018. Trust in Data Science: Collaboration, Translation, and Accountability in Corporate Data Science Projects. *Proceedings of the ACM on Human-Computer Interaction*. 2, CSCW (Nov. 2018), 1–28.
- [87] Passi, S. and Sengers, P. 2020. Making data science systems work. *Big data & society*. 7, 2 (2020), 205395172093960.
- [88] Pimentel, J.F., Murta, L., Braganholo, V. and Freire, J. 2019. A large-scale study about quality and reproducibility of jupyter notebooks. *Proceedings of the 16th International Conference on Mining Software Repositories (MSR)* (2019), 507–517.
- [89] Piorowski, D., González, D., Richards, J. and Houde, S. 2020. Towards evaluating and eliciting high-quality documentation for intelligent systems. *arXiv [cs.SE]*.
- [90] Piorowski, D., Park, S., Wang, A.Y., Wang, D., Muller, M. and Portnoy, F. 2021. How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study. *Proceedings of the ACM on Human-Computer Interaction*. 5.CSCW1 (2021), 1–25.
- [91] Polyzotis, N., Roy, S., Whang, S.E. and Zinkevich, M. 2018. Data Lifecycle Challenges in Production Machine Learning: A Survey. *ACM SIGMOD Record*. 47, 2

- (2018), 17–28.
- [92] Rahimi, M., Guo, J.L.C., Kokaly, S. and Chechik, M. 2019. Toward Requirements Specification for Machine-Learned Components. *Proceedings of the 27th International Requirements Engineering Conference Workshops (REW)* (2019), 241–244.
- [93] Rahman, M.S., Khomh, F., Hamidi, A., Cheng, J., Antoniol, G. and Washizaki, H. 2021. Machine Learning Application Development: Practitioners' Insights. *arXiv [cs.SE]*.
- [94] Rahman, M.S., Rivera, E., Khomh, F., Guéhéneuc, Y.-G. and Lehnert, B. 2019. Machine Learning Software Engineering in Practice: An Industrial Case Study. *arXiv [cs.SE]*.
- [95] Rakova, B., Yang, J., Cramer, H. and Chowdhury, R. 2020. Where Responsible AI meets Reality: Practitioner Perspectives on Enablers for shifting Organizational Practices. *Proceedings of the ACM on Human-Computer Interaction* (2020), 1–23.
- [96] Ribeiro, D.M., Cardoso, M., da Silva, F.Q.B. and França, C. 2014. Using qualitative metasummary to synthesize empirical findings in literature reviews. *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement* (2014), 1–4.
- [97] Ribeiro, M.T., Wu, T., Guestrin, C. and Singh, S. 2020. Beyond Accuracy: Behavioral Testing of NLP models with CheckList. *arXiv [cs.CL]*.
- [98] Riccio, V., Jahangirova, G., Stocco, A., Humbatova, N., Weiss, M. and Tonella, P. 2020. Testing machine learning based systems: a systematic mapping. *Empirical Software Engineering*, 25, 6 (2020), 5193–5254.
- [99] Rismani, S., Shelby, R., Smart, A., Jatho, E., Kroll, J., Moon, A. and Rostamzadeh, N. 2022. From plane crashes to algorithmic harm: applicability of safety engineering frameworks for responsible ML. *arXiv [cs.HC]*.
- [100] Riungu-Kalliosaari, L., Kauppinen, M. and Männistö, T. 2017. What Can Be Learnt from Experienced Data Scientists? A Case Study. *Product-Focused Software Process Improvement* (2017), 55–70.
- [101] Saha, D., Schumann, C., McElfresh, D.C., Dickerson, J.P., Mazurek, M.L. and Tschantz, M.C. 2020. Human Comprehension of Fairness in Machine Learning. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (2020), 152.
- [102] Salay, R., Queiroz, R. and Czarnecki, K. 2017. An Analysis of ISO 26262: Using Machine Learning Safely in Automotive Software. *arXiv [cs.AI]*.
- [103] Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P. and Aroyo, L.M. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), 1–15.
- [104] Sandelowski, M., Barroso, J. and Voils, C.I. 2007. Using qualitative metasummary to synthesize qualitative and quantitative descriptive findings. *Research in nursing & health*, 30, 1 (2007), 99–111.
- [105] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F. and Dennison, D. 2015. Hidden Technical Debt in Machine Learning Systems. *Advances in Neural Information Processing Systems* 28. C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, eds. Curran Associates, Inc. 2503–2511.
- [106] Sculley, D., Otey, M.E., Pohl, M., Spitznagel, B., Hainsworth, J. and Zhou, Y. 2011. Detecting adversarial advertisements in the wild. *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining* (2011), 274–282.
- [107] Sendak, M.P. et al. 2020. Real-World Integration of a Sepsis Deep Learning Technology Into Routine Clinical Care: Implementation Study. *JMIR medical informatics*, 8, 7 (2020), e15182.
- [108] Serban, A., van der Blom, K., Hoos, H. and Visser, J. 2020. Adoption and Effects of Software Engineering Best Practices in Machine Learning. *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)* (2020), 1–12.
- [109] Serban, A., van der Blom, K., Hoos, H. and Visser, J. 2021. Practices for Engineering Trustworthy Machine Learning Applications. *Proceedings of the 1st Workshop on AI Engineering - Software Engineering for AI (WAIN)* (2021), 97–100.
- [110] Serban, A. and Visser, J. 2022. Adapting Software Architectures to Machine Learning Challenges. *Proceedings of the 2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)* (2022), 152–163.
- [111] Shankar, S., Garcia, R., Hellerstein, J.M. and Parameswaran, A.G. 2022. Operationalizing Machine Learning: An Interview Study. *arXiv [cs.SE]*.
- [112] Shaw, M. and Zhu, L. 2022. Can Software Engineering Harness the Benefits of Advanced AI? *IEEE Software*, 39, 6 (2022), 99–104.
- [113] Siebert, J., Joekel, L., Heidrich, J., Nakamichi, K., Ohashi, K., Namba, I., Yamamoto, R. and Aoyama, M. 2020. Towards Guidelines for Assessing Qualities of Machine Learning Systems. *Proceedings of the International Conference on the Quality of Information and Communications Technology* (2020), 17–31.
- [114] Smith, D. 2017. Exploring development patterns in data science.
- [115] d. S. Nascimento, E., Ahmed, I., Oliveira, E., Palheta, M.P., Steinmacher, I. and Conte, T. 2019. Understanding Development Process of Machine Learning Systems: Challenges and Solutions. *Proceedings of the 2019 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)* (2019), 1–6.
- [116] Spencer, D. 2009. *Card Sorting: Designing Usable Categories*. Rosenfeld Media.
- [117] Studer, S., Bui, T.B., Drescher, C., Hanuschkin, A., Winkler, L., Peters, S. and Mueller, K.-R. 2021. Towards CRISP-ML(Q): A Machine Learning Process Model with Quality Assurance Methodology. *Machine Learning and Knowledge Extraction*, 3, 2 (2021), 392–413.
- [118] Tonekaboni, S., Joshi, S., McCradden, M.D. and Goldenberg, A. 2019. What Clinicians Want: Contextualizing Explainable Machine Learning for Clinical End Use. *Proceedings of the 4th Machine Learning for Healthcare Conference* (2019), 359–380.
- [119] Uchihira, N. 2022. Project FMEA for Recognizing Difficulties in Machine Learning Application System Development. *Proceedings of the 2022 Portland International Conference on Management of Engineering and Technology (PICMET)* (2022), 1–8.
- [120] Vogelsang, A. and Borg, M. 2019. Requirements Engineering for Machine Learning: Perspectives from Data Scientists. *Proceedings of the 27th International Requirements Engineering Conference Workshops (REW)* (2019), 245–251.
- [121] Wagstaff, K. 2012. Machine Learning that Matters. *arXiv [cs.LG]*.
- [122] Wang, D., Weisz, J.D., Muller, M., Ram, P., Geyer, W., Dugan, C., Tausczik, Y., Samulowitz, H. and Gray, A. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. *Proceedings of the ACM on Human-Computer Interaction*, 3, CSCW (2019), 1–24.
- [123] Wan, Z., Xia, X., Lo, D. and Murphy, G.C. 2019. How does Machine Learning Change Software Development Practices? *IEEE Transactions on Software Engineering*, 47, 9 (2019), 1857–1871.
- [124] Washizaki, H., Takeuchi, H., Khomh, F., Natori, N., Doi, T. and Okuda, S. 2020. Practitioners' insights on machine-learning software engineering design patterns: a preliminary study. *Proceedings of the 2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)* (2020), 797–799.
- [125] Washizaki, H., Uchida, H., Khomh, F. and Guéhéneuc, Y.-G. 2020. Machine learning architecture and design patterns. *IEEE Software*, 8, (2020).
- [126] Washizaki, H., Uchida, H., Khomh, F. and Guéhéneuc, Y.-G. 2019. Studying Software Engineering Patterns for Designing Machine Learning Systems. *Proceedings of the 10th International Workshop on Empirical Software Engineering in Practice (IWESEP)* (2019), 49–495.
- [127] Wohlin, C. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering. *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering* (2014), 1–10.
- [128] Zdanowska, S. and Taylor, A.S. 2022. A study of UX practitioners roles in designing real-world, enterprise ML systems. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022), 1–15.
- [129] Zhang, A.X., Muller, M. and Wang, D. 2020. How do data science workers collaborate? Roles, workflows, and tools. *Proceedings of the ACM on human-computer interaction*, 4, CSCW1 (2020), 1–23.
- [130] Zhang, J.M., Harman, M., Ma, L. and Liu, Y. 2022. Machine learning testing: Survey, landscapes and horizons. *IEEE Transactions on Software Engineering*, 48, 1 (2022), 1–36.
- [131] Zhang, X., Yang, Y., Feng, Y. and Chen, Z. 2019. Software Engineering Practice in the Development of Deep Learning Applications. *arXiv [cs.SE]*.
- [132] Zinkevich, M. 2017. Rules of machine learning: Best practices for ML engineering. minegrado.ovh.