



LUND UNIVERSITY

Exploring ML testing in practice - Lessons learned from an interactive rapid review with Axis Communications

Song, Qunying; Borg, Markus; Engström, Emelie; Ardö, Håkan ; Rico, Sergio

Published in:

2022 IEEE/ACM 1st International Conference on AI Engineering – Software Engineering for AI (CAIN)

2022

Document Version:

Early version, also known as pre-print

[Link to publication](#)

Citation for published version (APA):

Song, Q., Borg, M., Engström, E., Ardö, H., & Rico, S. (2022). Exploring ML testing in practice - Lessons learned from an interactive rapid review with Axis Communications. In *2022 IEEE/ACM 1st International Conference on AI Engineering – Software Engineering for AI (CAIN)* IEEE - Institute of Electrical and Electronics Engineers Inc.. <https://ieeexplore.ieee.org/document/9796447>

Total number of authors:

5

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Exploring ML testing in practice – Lessons learned from an interactive rapid review with Axis Communications

Qunying Song
Lund University
Lund, Sweden
qunying.song@cs.lth.se

Markus Borg
RISE Research Institutes of Sweden
and Lund University
Lund, Sweden
markus.borg@ri.se

Emelie Engström
Lund University
Lund, Sweden
emelie.engstrom@cs.lth.se

Håkan Ardö
Axis Communications
Lund, Sweden
hakan.x.arδο@axis.com

Sergio Rico
Lund University
Lund, Sweden
sergio.rico@cs.lth.se

ABSTRACT

There is a growing interest in industry and academia in machine learning (ML) testing. We believe that industry and academia need to learn together to produce rigorous and relevant knowledge. In this study, we initiate a collaboration between stakeholders from one case company, one research institute, and one university. To establish a common view of the problem domain, we applied an interactive rapid review of the state of the art. Four researchers from Lund University and RISE Research Institutes and four practitioners from Axis Communications reviewed a set of 180 primary studies on ML testing. We developed a taxonomy for the communication around ML testing challenges and results and identified a list of 12 review questions relevant for Axis Communications. The three most important questions (data testing, metrics for assessment, and test generation) were mapped to the literature, and an in-depth analysis of the 35 primary studies matching the most important question (data testing) was made. A final set of the five best matches were analysed and we reflect on the criteria for applicability and relevance for the industry. The taxonomies are helpful for communication but not final. Furthermore, there was no perfect match to the case company's investigated review question (data testing). However, we extracted relevant approaches from the five studies on a conceptual level to support later context-specific improvements. We found the interactive rapid review approach useful for triggering and aligning communication between the different stakeholders.

KEYWORDS

AI Engineering, Machine Learning, Testing, Interactive Rapid Review, Taxonomy

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

CAIN '22, May 22–23, 2022, Pittsburgh, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

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

ACM Reference Format:

Qunying Song, Markus Borg, Emelie Engström, Håkan Ardö, and Sergio Rico. 2022. Exploring ML testing in practice – Lessons learned from an interactive

rapid review with Axis Communications. In *Proceedings of International Conference on AI Engineering (CAIN '22)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Artificial intelligence (AI) applications have grown in popularity and pervasiveness. Among the AI applications currently in use, machine learning (ML) is the dominant technique with active communities in academia and industry [27]. Enterprises across diverse industry domains want to harness the new possibilities promoted by ML. However, due to their impact and increasing use in safety-critical domains, we need to develop ways to build trust in these applications. Bosch *et al.* calls for increased research on *AI engineering* [10], i.e., an evolution of software engineering practices and processes to meet the needs of systems development that incorporate trained ML models. These systems, in contrast to most traditional systems, have a probabilistic behavior [29]. Therefore, we need new approaches and solutions or adapt the existing solutions to new challenges [2, 8].

In this paper, we focus on ML testing, emphasizing applications of ML-based computer vision. ML testing has been a popular research topic in the last few years. Secondary studies show a rapidly increasing publication trend [33, 37, 46] and dedicated academic workshops and conferences have been established. As novel ML testing results are constantly published, both researchers and practitioners need ways to organize the information and sift through the massive academic output. Furthermore, there is a need for effective ways to match research proposals with application-specific industry needs [13].

An approach helpful in the inception of a collaborative project is interactive rapid reviews (IRRs) [34]. An IRR is a collaborative effort between researchers and practitioners that aims to identify and synthesize relevant research outcomes for the practitioners in their context. Apparently, an IRR could be a beneficial tool for new collaborative projects to explore interests, facilitate the exchange of ideas, and promote mutual understanding.

We conducted an IRR on ML testing with Axis Communications (hereafter Axis). The long-term goal of the IRR was to initiate a collaboration on ML testing between researchers from Lund University, RISE Research Institutes of Sweden, and practitioners at

Axis. As a means to that end, and a short-term goal, the IRR should identify the solution proposals from the academic community that are the most likely to provide value for Axis.

The contributions of this paper are the following. First, we developed a taxonomy about practical challenges and available research results on ML testing that helped us learn about the domain and navigate the research results (Sec. 4.1). Second, we compiled a list of twelve practical challenges, identified during the IRR at Axis, related to ML testing (Sec. 4.2). Third, we proposed a preliminary mapping, i.e., a potential connection between research results and practical challenges for three prioritized challenges at Axis (Sec. 4.3). Finally, we conducted an in-depth review of the 35 primary studies mapped to the highest priority topic, i.e., “How to test the dataset?” We extracted nine technological rules and identified context factors impacting the application of ML testing solutions found in the academic sources (Sec. 4.5).

2 BACKGROUND AND RELATED WORK

This section presents our position on AI quality and its connection to ML testing. Moreover, we introduce IRRs and the industrial case context.

2.1 AI quality and ML Testing

Quality is a multi-faceted concept that is notoriously difficult to nail down. Adding AI on top of this further exacerbates the challenge. Still, we posit that AI quality is going to be an increasingly important concept to ensure the trustworthiness that future AI systems must provide. AIQ is a regional effort to gather interested parties on AI quality, with a particular focus on the subset of AI that realizes functionality through supervised or unsupervised machine learning, i.e., MLware.

We adhere to the definition of AI quality as “the capability of MLware to satisfy stated and implied needs under specified conditions while the underlying data satisfy the requirements specific to the application and its context” [5]. The definition stresses that MLware combines data and conventional source code; thus, its quality is defined as an amalgamation of corresponding quality definitions from the IEC/ISO 25000 series [21, 22]. The definition is in line with discussions by Felderer *et al.* in the context of testing data-intensive systems [18]. Moreover, the emphasis on data quality assurance is central in this paper.

Inspired by Bjarnason *et al.*'s work on requirements engineering (RE) and software testing [3], our position is that AI quality assurance must be tackled from two directions. RE and testing must support MLware development projects as two bookends. As MLware is sensitive to changes, as Sculley *et al.* put it “changing anything changes everything” [36], aligning RE and testing is perhaps even more important than for conventional software engineering. Within AIQ, we have addressed RE for ML [7, 44], ML testing [9, 16, 30], and MLOps from the perspective of alignment [6]. In this paper, we again focus on ML testing.

ML testing is a rapidly growing research area that evolves software testing to meet the novel characteristics of ML-based systems. We used three secondary studies of ML testing [33, 37, 46] as the basis for the current work. The three secondary studies were considered the latest in the field when we initiated the current study. Given that the area ML testing is fairly new, we did not expect very

old publications. However, we did not explicitly exclude them either. The endpoint in the range was when we started the work, and we included all secondary studies we were aware of. Among the secondary studies we selected, two were published in 2020 [33, 46] and the other one in 2019 [37]. They used a systematic approach for searching, extracting, and synthesizing relevant studies on the topic of ML testing. It is also worth noting that Zhang *et al.* [46] and Ricco *et al.* [33] have also included arXiv pre-prints to be more extensive, and have identified 138 and 70 primary studies, respectively. In contrast, Sherin *et al.* restricted the literature search to peer-reviewed publications only and have identified 37 papers in their study [37]. In total, we have collected 180 unique primary studies based on the three secondary studies.

The three secondary studies report an increasing number of ML testing papers in recent years. The trend is the increased use of ML in various application domains and the importance of techniques for testing such applications. This is particularly evident as ML-based applications are deployed in both safety-critical and mission-critical contexts. Specifically, the majority of the studies that have been surveyed in Zhang *et al.* [46] are focusing on testing the correctness and robustness of supervised machine learning systems, while other type of learning such as unsupervised learning and reinforcement learning, and testing perspectives such as interpretability, efficiency, or privacy are much less studied. The analysis is consistent with the observations from Sherin *et al.* [37] that further attention is required to test the non-functional perspectives and different types of learning for ML systems. Sherin *et al.* [37] also highlighted that there is no adequate empirical evidence to evaluate the effectiveness of the available testing techniques, even though the area of ML testing keeps growing rapidly. In contrast, Ricco *et al.* [33] concluded that the most active research in ML testing has been dedicated on solving automatic test input generation and test oracle creation. Further studies are required to address numerous open challenges such as inventing proper testing metrics as well as benchmarks for ML systems.

2.2 Interactive rapid reviews

In this study, we use the guidelines for IRRs in software engineering proposed by Rico *et al.* [34]. Rapid reviews are a form of knowledge synthesis widely used in medicine to provide information quickly for decision-making. As an example, during the COVID-19 pandemic, a considerable amount of rapid reviews were conducted to support decision-making in many areas of medicine [41]. A group of researchers in software engineering proposed the use of rapid reviews to support decision-making in software projects [14]. A difference with the guidelines adopted in this study is the focus on the interaction between researchers and practitioners to make the reviews more relevant in the practitioners' context and foster the knowledge exchange.

The guidelines are presented as a series of steps. The first step is to *prepare the review*. In this step, an initial area and topic are identified, and the team is set. The second step is to *identify review questions and prepare the IRR protocol*, where based on the exchange, the IRR team formulate review questions that represent the common interest and plan the steps to conduct the IRR i.e. IRR protocol. Then, the third step is to *search and select papers*. In this

step, shortcuts are used to reduce the space of search. For that reason, it is suggested that the review questions are narrowed to specific questions. Then, during the fourth step, the IRR team *extracts and synthesizes data* from the research literature and prepare the actions to *disseminate IRR results* in the fifth step. It is important to clarify that the guidelines are flexible and may require adaption to the specific needs of the case.

The software engineering research community is familiar with systematic literature reviews (SLRs) as a form of knowledge synthesis. Although IRRs and SLRs are similar in many aspects, like methodology and the need for a systematic approach, it is important to clarify that IRRs do not pretend to be an alternative to SLRs when synthesising research literature. IRRs do not aim to be extensive, but provide rapid and valid input for the practitioners. IRRs address more narrow questions than SLRs. When conducting IRRs the review team applies shortcuts to narrow the search space and then save time in selecting and extracting relevant data. These shortcuts may result in missing relevant sources for the IRR. There are two main reasons to select an IRR for this study. Compared with SLRs, IRRs require less resources and can be completed in shorter time frames. Second, IRRs, as presented here, aim to promote exchange between researchers and practitioners, which is desirable at this phase of Axis.

2.3 Case context

Axis was the first industrial collaboration partner in AIQ. Within Axis, we identified a development team that develop solutions based on advanced ML-based computer vision. The team, develops people counting applications for dynamic environments such as shopping malls and public squares.

People counting is considered a “statistical application” that should be accurate on average. Corner cases are largely ignored, i.e., if a person wearing a “funny hat” is missed or double-counted is ok – as long as the counter is not incremented by an amount large enough to noticeably affect the hourly/daily statistics. This is in contrast to security surveillance applications for which corner cases are critical, e.g., possible intruders crawling under the camera. Still, accuracy over time is important to people counting. False positives (counting ghosts) and false negatives (missing people) are considered equally bad. Thus the F1-score (balanced harmonic mean of precision and recall) is the primary evaluation metric.

A set of test datasets representing scenarios in various operational environments is used for regression testing. As there are significant differences between operational environments, referred to as scenes, F1-scores are measured for individual scenes rather than for a single diverse test set. To provide reliable quality assurance, ensuring a high coverage of scenarios in the test dataset is essential. Differently sized regression test suites are running 1) in a continuous integration context, 2) on nightly builds, and 3) weekly. Two different test setups are used in the regression testing. One testing the algorithms involved only and one testing the actual hardware used.

3 METHOD

To initiate a collaboration on ML testing between researchers and practitioners at Lund University, RISE Research Institutes of Sweden, and Axis, we conducted an IRR [34] following the five steps in

Table 1. The expected outcome of the review was threefold: 1) to establish a common view of the general problem domain – ML testing, 2) to gain a quick overview of how current research matches with the specific needs at Axis, and 3) to propose a study aiming at filling one of the identified gaps. The researchers’ activities were carried out by the first three authors of this paper, while the fourth author and his colleagues represent the practitioner’s side. Finally, the fifth author guided and monitored the research procedure.

Table 1: Description of the five steps and corresponding research activities for this IRR study

Step	Activities for Researchers	Activities for Practitioners*
1. Prepare the review	Describe research area and preliminary research goals	Meeting to identify mutual information needs and agree on involvement.
2. Identify review questions and develop the IRR protocol	Propose SERP taxonomy, elaborate review questions and scope of search and selection.	Meeting to validate and refine taxonomy and elaborate on questions and scope. List and prioritize review questions.
3. Search and select papers	Map primary studies to review questions, iterate samples of selected studies with practitioners, update inclusion/exclusion criteria based on feedback.	Give feedback on relevance and applicability of selected papers.
4. Extract and synthesize data	Identify and assess maturity of technological rules.	Meeting to discuss results (esp. relevance, applicability) of technological rules. Discuss how and to whom results should be summarized and communicated.
5. Disseminate IRR results.	Design and present visual abstracts for the identified TRs. Propose new research studies.	Meeting to give feedback on results. Present results within company.

The steps of an IRR are similar to the steps of other types of systematic literature reviews but adapted to meet the specific needs of an industrial stakeholder. In our case the stakeholder was Axis.

3.1 Preparing the review

The goal of the preparation step was to form a review team of both researchers and practitioners and to identify mutual interests and information needs with respect to the general research topic, ML testing. In this step, the interaction between industry and academia took place in an input meeting. To prepare for the input meeting, the second author put together an overview presentation of the state-of-the-art of ML testing, and the fifth author put together an overview presentation of the IRR approach. Four researchers and four practitioners took part in the meeting. The four practitioners had different roles (expert engineer, software test engineer, senior software engineer, and technical leader) at the company and thus different perspectives on the topic. At the meeting, after the presentations, the practitioners shared aspects of their practices and challenges of testing their ML applications.

3.2 Identifying review questions and developing the review protocol

The goal of the second step was to agree on a list of prioritized review questions and an initial review protocol. Here interactions took place in a workshop and a follow-up ranking exercise. Before the workshop, the first three authors developed a preliminary SERP-taxonomy [31] based on the state-of-the-art of ML testing and previous SERP-taxonomies on software testing [1, 17]. A SERP-taxonomy includes four facets (i.e., scope, context, effect, and intervention) to align descriptions of research solutions and industry challenges. In that way, such a taxonomy may be used to facilitate communication between practitioners and support the mapping of challenges from the industry to available research [17].

During the workshop, including the full review team, we walked through all facets and entities of the taxonomy to trigger discussions about ML testing challenges and potential solutions from various perspectives. Based on the outcome of the workshop, the researchers updated the taxonomy and proposed a list of 12 potentially relevant review questions to Axis. This list was then sent out to all participants (researchers and practitioners) with a request to rank them in order of interest using an ordinal scale from 1 to 5. After summarizing the results of this exercise, we agreed to search for research relevant to the three highest ranked questions. Furthermore, we agreed to include research based on relevance and applicability for Axis, but did not specify this further at this point.

3.3 Searching and selecting primary studies

During this step, we successively refined the review protocol while searching for relevant studies and delimiting the scope (i.e., defined exclusion criteria). As part of this activity, the researchers conducted the search and selection while the practitioners gave feedback on relevance and applicability of a small sample of papers sent to them by email.

We limited the search to the research covered by three recent secondary studies on ML testing [33, 37, 46]. The researchers revisited the complete set of primary studies (180 papers) to map them to the three review questions. This screening was based on full-text scanning as it was impossible to do it based neither on the original classification (in the secondary studies) nor on a sole title and abstract screening. At this stage, we had an inclusive approach

meaning that if any of the three researchers marked a paper as relevant for a review question, it was coded as such.

Due to the large number of papers coded as potentially relevant to at least one of the three review questions, we decided to descope further and focus the in-depth analysis only on one of the review questions. Thus, the continued selection focused only on the highest prioritized review question, i.e., “How to test the dataset?” 35 of the 180 studies were marked as potentially relevant for this question. After a thorough review, only five of them remained. At this stage, the remaining candidates were tightly connected to the Axis’ context, i.e., testing data for ML-based computer vision.

3.4 Data extraction and synthesis

From the selected papers, we extracted technological rules following the design science lens described by Storey *et al.* [39]. A technological rule is a structured way of describing research contributions with respect to their effect, context, and intervention. Technological rules can be extracted and presented at different levels of abstraction, and are used for communicating the research output in a simple and condensed way [39]. Our goal was to compare technological rules, identify research gaps and the specific needs at Axis. To allow for generalization, we extracted technological rules of different abstraction levels from the primary studies. Furthermore, we extracted the maturity of the rules in terms of empirical observations and analytical reasoning that supported the propositions. The technological rules were then presented to the review team in a short reaction meeting, where the industrial team provided their reflections about relevance and applicability at Axis.

3.5 Disseminating the review results

As the main goal of the review was to initiate a new collaboration, we did not have a plan for disseminating the results within Axis. Instead the most relevant technological rules provide input for new MSc thesis proposals. However, unplanned dissemination took place as new knowledge travel between teams. The technological rules from one of the included papers were evaluated for another purpose by another team at Axis. On the academic side, results were summarized and presented as visual abstracts, one for each technological rule, at a seminar and in this report.

3.6 Validity of contributions

The main goal of conducting an IRR is not to build general theory or publish rigorous research results, but to extract sufficient knowledge to act in a specific situation. In our case, our primary goal was to align terminology, match interests, and initiate industry-academia collaboration on ML testing with one case company. Still, we believe the reported contributions could be helpful for other researchers with similar goals but they must be adapted to their contexts.

The taxonomy builds on previous work on taxonomy development [31], general testing terminology [17] and recent ML testing syntheses [33, 37, 46]. Hence, they are well founded in the research literature. However, the industrial validation is made from a narrow perspective in a single case context. Thus, when elaborating the taxonomy, the details mirror the review team’s interests and experiences, including four practitioners and four researchers.

Similarly, the list of challenges represents a single case, and the ranking of importance represents the interests of the review team. At a high abstraction level, these challenges and interests match the interests of the research community, represented by the 180 primary studies.

A review may also be validated based on its coverage, i.e., is any important work missing? In our case, we did not conduct an extensive search on our own but relied on the rigorous searches made in three recent secondary studies on the same topic. Although their searches were extensive, the time delay caused by the publication process leads to the omission of the most recent publications, i.e., from 2020 and onward. Thus, more recent research may provide a better match to our review questions. This should be considered in future work. Nevertheless, conclusions regarding relevance and applicability of available research are still valid and may guide future reviews.

4 RESULTS

In this section we present the outcome of each step of the IRR. Subsection 4.1 describes the entities of the taxonomy after validation, subsection 4.2 lists the open questions derived from the workshop, subsection 4.3 describes our mapping of primary studies to the review questions, subsection 4.4 presents the final exclusion criteria resulting from the iterative review of papers mapped to review question 1, subsection 4.5 presents the nine technological rules extracted from the five most relevant primary studies, and subsection 4.6 finally elaborates on what we did not find in the research literature, i.e., the gap between research and practice in our case.

4.1 ML testing taxonomy

As a result of the initial interaction between the case company and the researchers, we agreed on creating a taxonomy of ML testing to guide the collaboration further. We developed a general taxonomy for three out of four facets in the SERP-taxonomy architecture [31] (i.e., context, scope, and effect). Since we enter this review from the challenge perspective, we found that detailing these three facets were sufficient for the communication within the review team. The fourth SERP-facet, intervention, may be used to elaborate technical aspects that support abstraction and comparison of classes of solutions.

Our resulting taxonomy, presented in Figure 1, aims to guide the formulation of practical ML testing challenges at an appropriate abstraction level to support identification and design of relevant research.

4.1.1 Scope of ML testing interventions. The green sector in Figure 1 shows the details of the scope facet. Scope here refers to the testing activity, or part of the testing process, on which a potential intervention may focus. We identified four important aspects where two are derived from generic testing literature, i.e., testing process and testing levels, and the other two are specific to the ML context, i.e., parts of the ML-system to be tested and the mode of operations (online or offline testing, cf. Figure 1).

4.1.2 Effect of ML testing interventions. The effect of an intervention is described in terms of its observed or desired impact on the

testing. It could for example be the reported result of an empirical evaluation or an identified practical need in an exploratory study or case description. The blue sector in Figure 1 shows the effect facet. The main types of desired effects are derived from the generic SERP-taxonomy architecture [31], i.e., solving an unsolved problem (solve), improving the current solution (improve), and assessing the current situation (diagnose).

4.1.3 Context of ML testing interventions. The context facet aims to capture factors in the context that impact the effect or applicability of an intervention. Several such factors were identified, reflecting the multidimensional variation of both ML systems and testing approaches. Eight aspects of the context were included as shown in the yellow sector of Figure 1: 1) programming languages used for the implementation of the system, 2) degree of access to the system components, 3) framework, 4) machine learning type, 5) testing setup, 6) application, 7) type and 8) domain of the system.

At the initial stages of our project, we used the taxonomy for structuring the information in the secondary studies and for triggering discussions about the topic within the review team. We believe the resulting taxonomy may be used in similar ways by others to view ML testing challenges and solutions from various perspectives and thus support communication between researchers and practitioners who approach the topic – in different ways and in many cases at different abstraction levels. While conducting this review, we experienced that considering all the facets and digging in to details of all the facets of the taxonomy helped the communication and, by extension, our common understanding.

4.2 Prioritized list of open questions

Guided by the taxonomy and the feedback from practitioners in the initial meetings, we formulated 12 review questions potentially relevant to our case. In the following list, organized using the scope facet of the taxonomy in Figure 1, boxes represent the review team's highest priority questions. Verbs from the effect facet appear in italics.

- **ML system – Dataset**

- (#1) How to test the dataset? The dataset is evolving, which motivates the following sub-questions:
 - (a) How to identify mislabeled data in the dataset?
 - (b) Adequacy testing, i.e., how to assess and improve data (scenario) coverage of the training and test datasets in terms of diversity?
 - (c) How to assess potential bias after the training/test split?
 - (d) How valid is the data and its use? Is the data used for testing within the operational design domain? Or did some of the data originate from another source?

- **ML System – Learning program**

- (#2) How to *diagnose* (assess the fault detection capability) and *improve* (unit) testing (design and analysis) of the learning program?
- (#3) How to *improve* testing of the learning program to detect more faults? (e.g., using unit testing)

- **ML system – Learned model**

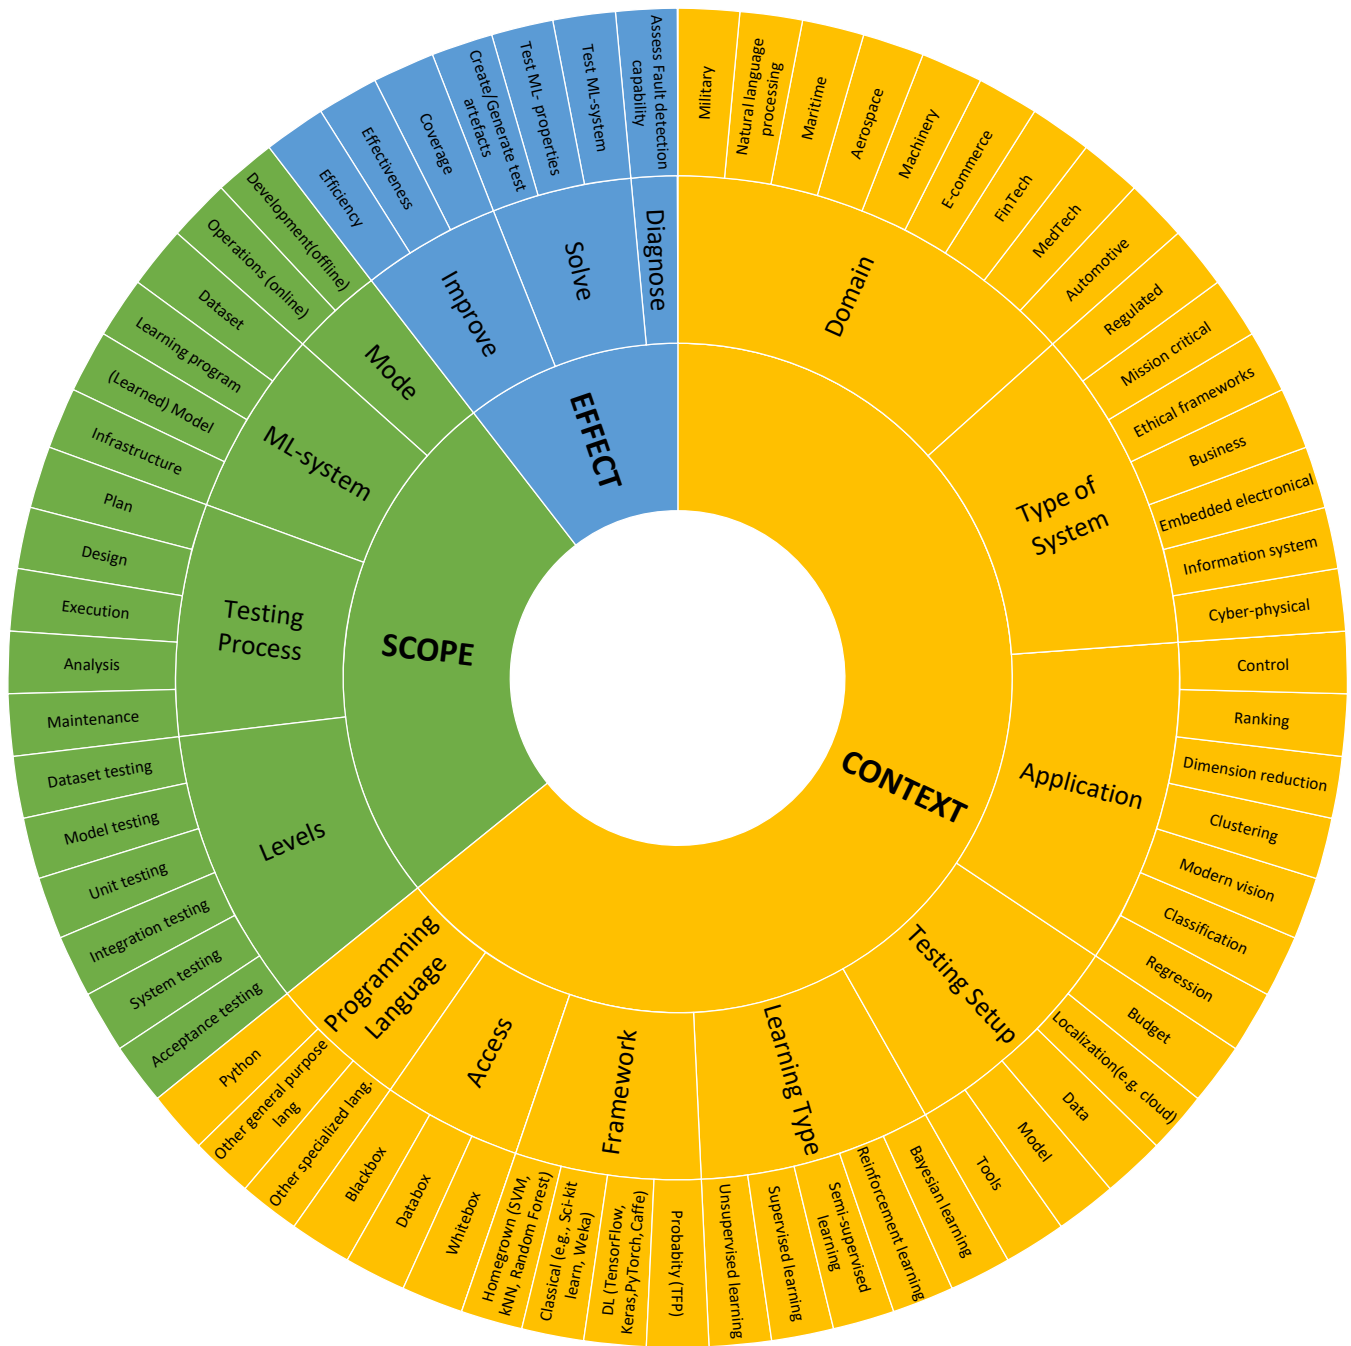


Figure 1: An ML testing SERP taxonomy with the facets context, scope, and effect.

(#4) Are there complementary metrics to assess model correctness (accuracy)? (e.g., edge case measures, uncertainty scores, aggregated metrics across scenes)

(#5) How to interpret and analyze the testing result of the ML model? (e.g., increased automation or visual analytics)

(#6) How to *diagnose* whether the current set of scenarios in the test dataset are appropriate for detecting faults on the model level?

(#7) How to *improve* coverage testing with respect to scenario diversity?

(#8) How to *improve* the test dataset to increase fault detection ability?

(#9) How to *improve* test prioritization to increase regression test efficiency?

(#10) How to generate new test cases for testing the model? (e.g. synthetic data, data augmentation, guided search)

• **Levels – System testing** (see Section 2.3)

(#11) How to *improve* (acceptance) testing of the ML-based system?

(#12) How to *diagnose* whether the current set of scenarios in the test dataset are appropriate for detecting faults on the system level?

Even though the above questions represent the interests of the review team, we argue that they constitute real challenges that generally deserve attention. We believe they may support other researchers trying to identify research gaps in ML testing for computer vision applications.

4.3 Mapping primary studies to the three most important open questions

The review team (i.e., four practitioners and three researchers) ranked the questions based on importance. Then the 180 primary studies covered in this review were mapped to the three review questions that were considered the most important (i.e., number #1, #4, and #10 (in the list of questions above)). 35 primary studies were marked as potentially relevant for review question 1, 25 primary studies for review question 4, and 68 for review question 10. The details of this mapping can be found on Zenodo [38].

This mapping may help navigating the research literature and could be used as a starting point by researchers and practitioners facing similar challenges.

4.4 Selection of studies related to question #1

This section describes our analysis of relevance and applicability for Axis in relation to the review question #1. The final list of exclusion criteria is:

- **Purpose.** The proposed mechanism does not evaluate some properties of the data. In MLware, the training dataset is part of the system [5]. The purpose of the proposed mechanism shall be to test these data, not to generate synthetic data – unless the synthetic data are specifically used to validate the training data through comparison. Examples of excluded papers relate to:
 - Test data generation for ML systems such as Zhang *et al.* [47] and Tian *et al.* [40] that generate artificial driving scenes for testing ML-based autonomous driving functions.
 - Testing the model fitness like Zhang *et al.* [45] which validates the model relevance, and ML model underfitting as well as overfitting using a perturbed model validation technique.
 - Online monitoring of data prior to making predictions, e.g., Henriksson *et al.* [19], since it targets testing aspects of the operational environment and detects inputs that are outside the training dataset.
 - The proposed mechanism is intended to protect the neural network from antagonistic attacks. Antagonistic attacks

are closer to cybersecurity research than data validation. Furthermore, adversarial attacks would typically target the system in operation and not the training/validation/test dataset. An example of such studies is Uesato *et al.* [42] which evaluates learning systems in safety-critical domains by identifying the adversarial situations.

- **Applicability.** The proposed mechanism is not applicable in the Axis' context. Reasons for exclusion include:
 - *Not NN learning.* The mechanism is not applicable to supervised learning with neural networks. Axis uses neural networks for supervised learning and the proposed mechanism must be applicable. Thus, papers that explicitly address other learning mechanisms are considered out of scope, such as Krishnan *et al.* [26] on support vector machines and Uesato *et al.* [42] on reinforcement learning.
 - *Not images.* There is no explicit mention of how the proposed mechanism could work for images. Axis trains models for video sequences and the proposed mechanism must be applicable. Thus, interventions that validate only non-image data are excluded, e.g., spell or format checking, and named entity recognition on tabular data [11, 20, 35].

4.5 Best matches

After applying the criteria described above, five primary studies remained, partly answering the general review question (#1 in Section 4.2) “How to test the dataset?” Although none of the proposed solutions were directly applicable in the case context, Axis confirmed related problem formulations and shared potentially valuable ideas. In line with the design science lens [39], we extracted technological rules for each paper. A technological rule captures the mapping between a problem and a solution. We describe them in terms of: **To achieve <effect> IN <context> DO <intervention>**. We describe the technological rules at different abstraction levels, i.e., some are more concrete and others are more general.

4.5.1 Paper 1. Ma *et al.* [28] propose a mutation testing framework for deep learning systems to assess the quality of the test data. Specifically, a set of source-level mutation operators are defined to introduce faults to the training data and the training programs. In addition, a set of model-level mutation operators are defined to create mutants for the deep learning models without a training process. Eventually, the effectiveness of the test data can be evaluated from the analysis based on to what extent the injected faults could be detected. The authors have used the framework on two publicly available datasets (i.e., MNIST and CIFAR-10) with three popular deep learning models, and demonstrated the effectiveness of the framework for designing and constructing high-quality test datasets for deep learning MLware.

Two technological rules were extracted from this paper [28], i.e., a concrete one:

Technological rule 1

To measure the quality of test data for deep learning systems, use an adapted form of mutation testing.

and a general one:

Technological rule 2

To improve the generality and robustness of deep learning models, test the test dataset.

Both the source-level and model-level mutant operators are general and can be reused. Furthermore, the concept of the proposed framework, i.e., to use mutation testing, is not constrained by any specific type of deep learning application or data used. Thus, we consider this paper potentially relevant, and the synthesized findings can be transformed for testing the quality of test data for Axis as well. The response from Axis was positive, they found the approach interesting but questioned the scalability. Axis works with orders of magnitude larger datasets (about 10^5 to 10^7 images, which is 10 to 1,000 times larger than the MNIST and CIFAR-10 datasets) than the ones used for evaluating the approach in the paper. In addition, the industrial datasets are rather Full HD resolution than the 32×32 pixels targeted in many research papers. Another question is which mutant operators would work for the complexity of the industrial case, as MNIST and CIFAR-10 are trivial datasets in comparison. A pre-study in the Axis context will be initiated to investigate this further.

4.5.2 Paper II. Kim *et al.* [24] propose a concept of surprise adequacy as the test adequacy criterion for deep learning systems. Based on the trace of the neuron activation when executing the deep learning model on both the training data and the testing data, surprise adequacy can be calculated using either a likelihood-based approach or a distance-based approach. The resulting surprise adequacy indicates how surprising the test data is compared to the training data. The concept assumes that a good test input should be sufficiently, but not excessively, surprising to the training data. The authors also evaluate the effectiveness of using the surprise adequacy metric for sampling test input and improving the model accuracy via retraining, based on publicly available datasets such as MNIST and CIFAR-10, and deep learning systems for autonomous vehicles like Dave-2 and Chauffeur.

Two technological rules were extracted from this paper [24], the concrete one is:

Technological rule 3

To improve the classification accuracy of deep learning systems, retrain the model, systematically sampling inputs based on the surprise adequacy criterion.

while the general one is described as:

Technological rule 4

To test the correctness and robustness of deep learning systems, test the systems' behaviour with respect to their training data.

The proposed criterion – surprise adequacy – and the corresponding ways of computing such a criterion are transferable to different deep learning MLware as the training data, testing data, and neuron activations are inherent parts of such systems. The

expected outcome is an indication of how good or how different the test is compared to the training data.

While the surprise adequacy metric was not considered relevant for the practitioners in the review team, it was explored by another team at Axis. Here it was explored for a slightly different purpose, i.e., to guide complementary data collection rather than for testing data. Collecting additional data that strives to maximize the surprise adequacy has also been proposed by the original inventors [25], as an approach to increase the diversity of both training and test datasets. However, Axis compared surprise adequacy to a set of other metrics and in the end they selected another option (recent work proposed by Pleiss *et al.* [32], not covered by the secondary studies used for our study selection). Part of the reason was that the proposed surprise adequacy calculations did not scale to size of the data set as described earlier (see subsection 4.5.1). On the other hand, Axis encourages additional research into surprise adequacy calculation for representative subsets of the data, i.e., aggregating measures for families of neuron activation traces.

4.5.3 Paper III. Byun *et al.* [12] introduce three different metrics for test input prioritization for deep neural networks, i.e., (1) confidence, measured by using the softmax function, (2) uncertainty, measured using Bayesian Networks, and (3) surprise as described in the previous paper by Kim *et al.* [24]. The authors apply these metrics to prioritize test inputs for two different systems (i.e., a digital classification system trained on the MNIST dataset, and TaxiNet) for image classification. They show the effectiveness of the metrics in indicating fault-revealing inputs and, by extension, for selecting test input and improving the model via retraining.

Two technological rules were extracted from this paper [12], including a concrete one:

Technological rule 5

To prioritize test input for deep neural networks, apply metrics of confidence, uncertainty, and surprise.

and a general one:

Technological rule 6

To increase test effectiveness when testing deep neural networks in safety-critical systems, prioritize test input.

Similar to paper II, this paper proposes metrics to measure the sentiment of the deep neural networks. Then, the test inputs can be validated, prioritized, and effectively selected for testing and retraining the model. The findings are considered generic and potentially relevant for Axis' needs. While "Surprise Adequacy" was investigated by another team at the company (see subsection 4.5.2) the other two metrics "confidence" and "uncertainty" have not yet been considered.

4.5.4 Paper IV. Cheng *et al.* [15] study a set of metrics to measure the dependability attributes of neural networks. The metrics include robustness, interpretability, completeness, and correctness. In our review, the paper was initially excluded due to its purpose (i.e., measuring the dependability of neural network models) and the application on autonomous driving. However, after a second

review, the paper was included since the part related to the completeness of the training data seems relevant. The paper uses neuron k-activation and neuron activation patterns as a measure of scenario coverage and completeness of training data for NN-based autonomous driving systems. The assessment of completeness of the training data, which is used for testing the coverage of the training data but is general for testing of data regardless of its use, could support both quality assurance of training or test datasets.

One general technological rule was extracted from this paper [15] as only the part about the training data completeness is relevant. The technological rule is described as:

Technological rule 7

To measure dependability of neural networks, evaluate the completeness of the training data.

The part that involves measuring completeness of training data is relevant for data testing in general. While the paper sets the general focus in autonomous driving applications, further investigation on how it could be applied into Axis' context should be performed.

4.5.5 *Paper V. Bolte et al.* [4] construct a system framework for corner case detection in training datasets for autonomous driving systems. The framework consists of three parts: (1) a semantic segmentation model to partition and classify the image into different semantic parts; (2) an image prediction model that predicts the next image based on the previous set of images and counting the errors based on the real image; and (3) a detection system that detects the corner cases if an object is unpredictable given the error score counted in the previous model. The proposed framework can be used in both online and offline modes. The difference is that the offline mode takes a collected database of image data for training, and the online model uses live video frames collected by the camera installed on vehicles. The authors have trained and evaluated the framework on the Cityscapes dataset and achieved prominent results for detecting unusual situations for autonomous driving. Two technological rules were extracted from this paper [4], the concrete one is:

Technological rule 8

To detect corner cases for ML-based autonomous driving systems, use a system framework based on image segmentation and prediction.

while the general one is:

Technological rule 9

To improve the robustness of machine learning systems, identify critical situations.

The concept of the study is relevant for data testing for ML systems in general, although the proposed system framework works with image data. It predicts corner cases for autonomous driving, which can be used to test and retrain the ML model. Still, the actual applicability of using this framework in Axis needs to be further investigated.

While identifying corner cases is not so important for the application of people counting, which is the focus of the practitioners in the review team, it could be vital for the companies' other products in the security business segment. For example, a corner case of someone moving strangely to avoid detection would be critically important for a security application to detect. For security applications, detecting such anomalies are among the most important use cases. On the other hand, for people counting applications, it could be interesting to apply a broader definition of corner cases, i.e., not only very rare cases but rather underrepresented scenarios. Any mechanism that could support identification of such scenarios would be helpful for ML testing.

4.6 Identified gap

The five papers that we have presented in the previous subsection mainly include frameworks, metrics, constructed tools, and mechanisms for measuring the quality of the data and consistency of the test data with respect to the training data for ML systems. Based on the general focus and sub-questions of the first review question (i.e., how to test the dataset), we observed that the extracted technological rules can be used to address some, but not all, related perspectives.

In particular, there are no studies that provide relevant techniques for a) identifying mislabeled data in the dataset. Thus, this sub-question is still an open challenge and needs to be studied further. However, paper IV provides mechanisms for measuring the completeness of training data, which gives some insights and potential solutions for b) assessing and improving the training and testing data coverage – also from the perspective of scenario coverage. Note that the proposed mechanism focuses on the autonomous driving domain. In addition, no studies we identified target c) assessing biases between the training and testing data. However, papers II and III could be potentially relevant since they support measuring and filtering test data not represented in the training data using different metrics (i.e., confidence, uncertainty, and surprise adequacy). The same observation also applies for sub-question d) how valid the data used for testing is with respect to the training data, where the metrics proposed in papers II and III can be used for such purposes. As a result, we believe the findings we synthesized can ease some parts of the research question we focus on, whereas further investigation is needed to address the remaining gaps.

5 DISCUSSION

AI engineering is an emerging field that is vital for AI quality. As argued in Section 2, ML testing is going to play a critical role in ensuring that future AI solutions are trustworthy. However, there is no established go-to model describing ML testing. Several previous studies propose dimensions to bring structure to the research area. This paper synthesizes a novel taxonomy based on three secondary studies. The taxonomy shall be considered work-in-progress, but it already has provided value for us in an emerging industry-academia collaboration.

The three secondary studies used three different classification strategies for their respective goals. Both Riccio *et al.* and Sherin *et al.* refer to their works as systematic mapping studies and focus on trends and gaps. The scope of Riccio *et al.* is functional testing

of ML systems and they structure the 70 primary studies based on 1) system context, 2) testing approach, and 3) empirical evaluation [33]. Sherin *et al.*'s mapping, including also non-functional testing, provides less synthesis and rather extracts fine-granular information from the 37 primary studies. In the secondary study by Zhang *et al.*, referred to only as a survey, the authors explicitly specify their ambition to provide a comprehensive overview of ML testing. The 138 included papers are organized into 1) testing properties, 2) testing components, 3) testing workflow, and 4) application scenarios. We believe that our proposed taxonomy combines the complementary perspectives provided in previous work.

Any taxonomy or model is developed for a specific communication purpose, targeting a defined group of people. In our case, the goal was to align terminology and identify shared interests within the group of researchers and practitioners. Thus, we found the SERP approach [31] applicable and useful. Furthermore, by building on SERP-test [17], comprising common testing terminology, and adding the ML perspective from the secondary studies [33, 37, 46] as well as from the case company, we got a solid basis for our communication on the topic – ML testing.

At a general level, we found a good match between the prioritized challenges of our case company and the research focus within the community. As a result, 92 out of 180 primary studies were classified as potentially relevant for at least one of the top three review questions, i.e., #1 How to test the dataset (35 primary studies), #2 How to assess model accuracy (metrics) (25 primary studies), and #3 How to generate test cases for testing the model (68 primary studies). However, as shown in the in-depth analysis of the first set of papers, no perfect matches exist.

Of the 35 primary studies initially considered relevant for data testing, we finally selected and reviewed five that best matched the specific context and needs at Axis. After analyzing and synthesizing the findings from the papers, we extracted nine technological rules. The practitioners were positive to the presented techniques (see subsection 4.5) and thought most of them could be relevant. Particularly, they have used the surprise adequacy for complementing data collection in the company as described in subsection 4.5.2. However, we found no perfect match directly transferable to the applications and data testing issues at Axis. As underlined by the definition of AI quality [5], finding feasible ways to perform data quality assurance is at the heart of the problem. Therefore, we are convinced that data testing will play an important role in the future of AI engineering. Also, it is significant in future research to explore further how to instantiate and evaluate the techniques in this industrial setting.

The concepts (e.g., metrics and criteria) and interventions (e.g., approaches and frameworks) presented in the five papers are quite generic for data testing in the ML field as to the extent of our interpretation. Hence, we believe those concepts and interventions can be reused and adapted for solving potential issues for different ML application domains. In the same way, the technological rules can be used to map solutions to challenges at different abstract levels, and support the communication and knowledge exchange between academic researchers and industrial practitioners in further studies.

6 CONCLUSIONS

We report outcomes and lessons learned from applying an interactive rapid review on machine learning testing. The review team

consisted of four researchers from Lund University and RISE Research Institutes of Sweden and four practitioners from Axis. The primary goal of the study was to initiate collaboration and align terminology and interests between the partners.

Three secondary studies, covering 180 primary studies on machine learning testing, functioned as a starting point for the review. The classifications of research in the secondary studies were mapped to the SERP taxonomy architecture [31] to guide the alignment of terminology and interests within the review team. The resulting SERP taxonomy were further extended by general taxonomies on software testing [17] built on the same taxonomy architecture. Finally, we validated and updated the outcome based on discussions and reflections in the review team. While we plan to evolve the outcome, *this paper presents the latest version of the taxonomy.*

The new SERP taxonomy was used to identify and describe current challenges in the case context. *The complete list of challenges are presented in this report.* Furthermore, the review team ranked the challenges by their perceived importance for the target organizational unit within the case company.

The primary studies were mapped to the three most important questions. Moreover, we conducted an in-depth analysis of the 35 papers for the highest ranked question, i.e., “How to test the dataset?” *We present and discuss 9 technological rules on data testing, extracted from 5 of the papers. Finally, we report and discuss the relevance and applicability criteria used to filter out those 5 papers.*

As AI quality combines source code and data quality [5], we believe that data testing will be increasingly important within the field of AI engineering. Our findings call for more research on the topic, not the least for image data, targeting business-critical computer vision systems. Furthermore, convincing data testing for computer vision applications can potentially constitute a cornerstone in the safety argumentation in future assurance cases, e.g., for critical ML-based perception applications in automotive [8], avionics [43], and healthcare [23].

The motivation for this interactive rapid review was to identify research or research gaps of relevance for the case company. Thus, all steps in the process have been guided by Axis' specific needs. As our next step, we plan to design a joint solution-oriented study on the topic of data testing as well as a set of MSc thesis project proposals. Based on the discussions of the selected studied in Section 4.5, there are several promising directions for future collaborations. Our case (of industry-academia collaboration) is a single case and as such a proof-of-concept that may be extended with additional cases.

ACKNOWLEDGMENTS

This initiative received financial support through the AIQ Meta-Testbed project funded by Kompetensfonden at Campus Helsingborg, Lund University, Sweden. In addition, this work was supported in part by the Wallenberg AI, Autonomous Systems and Software Program (WASP).

REFERENCES

- [1] Nauman Bin Ali, Emelie Engström, Masoumeh Taromirad, Mohammad Mousavi, Nasir Mehmood Minhas, Daniel Helgesson, Sebastian Kunze, and Mahsa Varshoaz. 2019. On the search for industry-relevant regression testing research. *Empirical*

- Software Engineering* 24, 4 (2019), 2020–2055. <https://doi.org/10.1007/s10664-018-9670-1>
- [2] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. 2019. Software Engineering for Machine Learning: A Case Study. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. 291–300. <https://doi.org/10.1109/ICSE-SEIP.2019.00042>
 - [3] Elizabeth Bjarnason, Per Runeson, Markus Borg, Michael Unterkalmsteiner, Emelie Engström, Björn Regnell, Giedre Sabaliauskaite, Annabella Loconsole, Tony Gorschek, and Robert Feldt. 2014. Challenges and practices in aligning requirements with verification and validation: a case study of six companies. *Empirical software engineering* 19, 6 (2014), 1809–1855.
 - [4] Jan-Aike Bolte, Andreas Bar, Daniel Lipinski, and Tim Fingscheidt. 2019. Towards corner case detection for autonomous driving. In *2019 IEEE Intelligent vehicles symposium (IV)*. IEEE, 438–445.
 - [5] Markus Borg. 2021. The AIQ Meta-Testbed: Pragmatically Bridging Academic AI Testing and Industrial Q Needs. In *Software Quality: Future Perspectives on Software Engineering Quality*. Dietmar Winkler, Stefan Biffl, Daniel Mendez, Manuel Wimmer, and Johannes Bergsmann (Eds.). Springer International Publishing, Cham, 66–77.
 - [6] Markus Borg. 2022. Agility in Software 2.0 – Notebook Interfaces and MLOps with Buttresses and Rebars. In *Proc. of the International Conference on Lean and Agile Software Development*. Springer, 3–16.
 - [7] Markus Borg, Joshua Bronson, Linus Christensson, Fredrik Olsson, Olof Lennartsson, Elias Sonnsjö, Hamid Ebadi, and Martin Karsberg. 2021. Exploring the Assessment List for Trustworthy AI in the Context of Advanced Driver-Assistance Systems. In *2021 IEEE/ACM 2nd International Workshop on Ethics in Software Engineering Research and Practice (SEthics)*. IEEE, 5–12.
 - [8] Markus Borg, Cristofer Englund, Krzysztof Wnuk, Boris Duran, Christoffer Lewandowski, Shenjian Gao, Yanwen Tan, Henrik Kaijser, Henrik Lönn, and Jonnas Törnqvist. 2019. 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 (2019), 1–19.
 - [9] Markus Borg, Ronald Jabangwe, Simon Åberg, Arvid Ekblom, Ludwig Hedlund, and August Lidfeldt. 2021. Test automation with grad-CAM Heatmaps-A future pipe segment in MLOps for Vision AI?. In *2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 175–181.
 - [10] Jan Bosch, Helena Holmström Olsson, and Ivica Crnkovic. 2021. Engineering ai systems: A research agenda. In *Artificial Intelligence Paradigms for Smart Cyber-Physical Systems*. IGI Global, 1–19.
 - [11] Eric Breck, Neoklis Polyzotis, Sudip Roy, Steven Whang, and Martin Zinkevich. 2019. Data Validation for Machine Learning. In *MLSys*.
 - [12] Taejoon Byun, Vaibhav Sharma, Abhishek Vijayakumar, Sanjai Rayadurgam, and Darren Cofer. 2019. Input prioritization for testing neural networks. In *2019 IEEE International Conference on Artificial Intelligence Testing (AITest)*. IEEE, 63–70.
 - [13] Anita D Carleton, Erin Harper, Michael R Lyu, Sigrid Eldh, Tao Xie, and Tim Menzies. 2020. Expert Perspectives on AI. *IEEE Software* 37, 4 (2020), 87–94.
 - [14] Bruno Cartaxo, Gustavo Pinto, and Sergio Soares. 2020. *Rapid Reviews in Software Engineering*. Springer International Publishing, Cham, 357–384. https://doi.org/10.1007/978-3-030-32489-6_13
 - [15] Chih-Hong Cheng, Chung-Hao Huang, Harald Ruess, Hirotohi Yasuoka, et al. 2018. Towards dependability metrics for neural networks. In *2018 16th ACM/IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)*. IEEE, 1–4.
 - [16] Hamid Ebadi, Mahshid Helali Moghadam, Markus Borg, Gregory Gay, Afonso Fontes, and Kasper Socha. 2021. Efficient and Effective Generation of Test Cases for Pedestrian Detection-Search-based Software Testing of Baidu Apollo in SVL. In *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*. IEEE, 103–110.
 - [17] Emelie Engström, Kai Petersen, Nauman Bin Ali, and Elizabeth Bjarnason. 2017. SERP-Test: A Taxonomy for Supporting Industry—Academia Communication. *Software Quality Journal* 25, 4 (dec 2017), 1269–1305. <https://doi.org/10.1007/s11219-016-9322-x>
 - [18] Michael Felderer, Barbara Russo, and Florian Auer. 2019. On testing data-intensive software systems. In *Security and Quality in Cyber-Physical Systems Engineering*. Springer, 129–148.
 - [19] Jens Henriksson, Christian Berger, Markus Borg, Lars Tornberg, Cristofer Englund, Sankar Raman Sathayamoorthy, and Stig Ursing. 2019. Towards structured evaluation of deep neural network supervisors. In *2019 IEEE International Conference on Artificial Intelligence Testing (AITest)*. IEEE, 27–34.
 - [20] Nick Hynes, D Sculley, and Michael Terry. 2017. The data linter: Lightweight, automated sanity checking for ml data sets. In *NIPS MLSys Workshop*.
 - [21] ISO/IEC. 2008. ISO 25012 Systems and software engineering – Systems and software quality requirements and evaluation (SQuaRE) - Data quality model.
 - [22] ISO/IEC. 2011. ISO 25010 Systems and software engineering – Systems and software quality requirements and evaluation (SQuaRE) - System and software quality models.
 - [23] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: past, present and future. *Stroke and vascular neurology* 2, 4 (2017).
 - [24] Jinhwan Kim, Robert Feldt, and Shin Yoo. 2019. Guiding deep learning system testing using surprise adequacy. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 1039–1049.
 - [25] Jinhwan Kim, Jeongil Ju, Robert Feldt, and Shin Yoo. 2020. Reducing dnn labelling cost using surprise adequacy: An industrial case study for autonomous driving. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1466–1476.
 - [26] Sanjay Krishnan, Jiannan Wang, Eugene Wu, Michael J. Franklin, and Ken Goldberg. 2016. ActiveClean: Interactive Data Cleaning for Statistical Modeling. *Proc. VLDB Endow.* 9, 12 (aug 2016), 948–959.
 - [27] Lucy Ellen Lwakatere, Aiswarya Raj, Ivica Crnkovic, Jan Bosch, and Helena Holmström Olsson. 2020. Large-scale machine learning systems in real-world industrial settings: A review of challenges and solutions. *Information and Software Technology* 127 (2020), 106368. <https://doi.org/10.1016/j.infsof.2020.106368>
 - [28] Lei Ma, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Felix Juefei-Xu, Chao Xie, Li Li, Yang Liu, Jianjun Zhao, et al. 2018. Deepmutation: Mutation testing of deep learning systems. In *2018 IEEE 29th International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 100–111.
 - [29] Dusica Marijan, Arnaud Gotlieb, and Mohit Kumar Ahuja. 2019. Challenges of Testing Machine Learning Based Systems. In *2019 IEEE International Conference on Artificial Intelligence Testing (AITest)*. 101–102. <https://doi.org/10.1109/AITest.2019.00010>
 - [30] Mahshid Helali Moghadam, Markus Borg, and Seyed Jalaleddin Mousavirad. 2021. Deeper at the sbst 2021 tool competition: ADAS testing using multi-objective search. In *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*. IEEE, 40–41.
 - [31] Kai Petersen and Emelie Engström. 2014. Finding Relevant Research Solutions for Practical Problems: The Serp Taxonomy Architecture. In *Proceedings of the 2014 International Workshop on Long-Term Industrial Collaboration on Software Engineering (WISE '14)*. Association for Computing Machinery, New York, NY, USA, 13–20. <https://doi.org/10.1145/2647648.2647650>
 - [32] Geoff Pleiss, Tianyi Zhang, Ethan Elenberg, and Kilian Q Weinberger. 2020. Identifying Mislabeled Data using the Area Under the Margin Ranking. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. 17044–17056.
 - [33] Vincenzo Riccio, Gunel Jahangirova, Andrea Stocco, Nargiz Humbatova, Michael Weiss, and Paolo Tonella. 2020. Testing machine learning based systems: a systematic mapping. *Empirical Software Engineering* 25, 6 (2020), 5193–5254.
 - [34] Sergio Rico, N. Ali, Emelie Engström, and Martin Höst. 2020. Guidelines for conducting interactive rapid reviews in software engineering – from a focus on technology transfer to knowledge exchange.
 - [35] Sebastian Schelter, Dustin Lange, Philipp Schmidt, Meltem Celikel, Felix Biessmann, and Andreas Grafberger. 2018. Automating large-scale data quality verification. *Proceedings of the VLDB Endowment* 11, 12 (2018), 1781–1794.
 - [36] D. Sculley et al. 2015. Hidden Technical Debt in Machine Learning Systems. In *Proc. of the 28th Int'l Conf. on Neural Information Proc. Systems*. 2503–2511.
 - [37] Salman Sherin, Muhammad Uzair Khan, and Muhammad Zohaib Iqbal. 2019. A Systematic Mapping Study on Testing of Machine Learning Programs. arXiv:cs.LG/1907.09427
 - [38] Qunying Song, Markus Borg, Emelie Engström, Håkan Ardö, and Sergio Rico. 2022. Primary Studies. <https://doi.org/10.5281/zenodo.5865070>
 - [39] Margaret-Anne Storey, Emelie Engstrom, Martin Höst, Per Runeson, and Elizabeth Bjarnason. 2017. Using a Visual Abstract as a Lens for Communicating and Promoting Design Science Research in Software Engineering. In *2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. 181–186. <https://doi.org/10.1109/ESEM.2017.28>
 - [40] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th international conference on software engineering*. 303–314.
 - [41] Andrea C Tricco, Chantelle M Garritty, Leah Boulos, Craig Lockwood, Michael Wilson, Jessie McGowan, Michael McCaul, Brian Hutton, Fiona Clement, Nicole Mittmann, et al. 2020. Rapid review methods more challenging during COVID-19: commentary with a focus on 8 knowledge synthesis steps. *Journal of clinical epidemiology* 126 (2020), 177–183.
 - [42] Jonathan Uesato, Ananya Kumar, Csaba Szepesvari, Tom Erez, Avraham Ruderman, Keith Anderson, Nicolas Heess, Pushmeet Kohli, et al. 2018. Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures. *arXiv preprint arXiv:1812.01647* (2018).
 - [43] Guillaume Vidot, Christophe Gabreau, Ileana Ober, and Iulian Ober. 2021. Certification of embedded systems based on Machine Learning: A survey. *arXiv preprint arXiv:2106.07221* (2021).
 - [44] Andreas Vogelsang and Markus Borg. 2019. Requirements Engineering for Machine Learning: Perspectives from Data Scientists. In *Proc. of the 27th Int'l Requirements Engineering Conf. Workshops*. 245–251. <https://doi.org/10.1109/REW.2019.00050>

- [45] Jie Zhang, Earl Barr, Benjamin Guedj, Mark Harman, and John Shawe-Taylor. 2019. Perturbed model validation: A new framework to validate model relevance. (2019).
- [46] Jie M Zhang, Mark Harman, Lei Ma, and Yang Liu. 2020. Machine learning testing: Survey, landscapes and horizons. *IEEE Transactions on Software Engineering* (2020).
- [47] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. 2018. DeepRoad: GAN-based metamorphic testing and input validation framework for autonomous driving systems. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 132–142.