



CHALMERS
UNIVERSITY OF TECHNOLOGY



UNIVERSITY OF GOTHENBURG

A Domain-Specific Language for Cross-platform, Edge-deployed Machine Learning Models

A Model Interpretation-based Approach

Master's thesis in Computer science and engineering

Albin Karlsson Landgren
Philip Perhult Johnsen

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden 2024

MASTER'S THESIS 2024

A Domain-Specific Language for Cross-platform, Edge-deployed Machine Learning Models

A Model Interpretation-based Approach

Albin Karlsson Landgren
Philip Perhult Johnsen



UNIVERSITY OF
GOTHENBURG



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden 2024

A Domain-Specific Language for Cross-platform, Edge-deployed Machine Learning
Models

A Model Interpretation-based Approach

ALBIN KARLSSON LANDGREN

PHILIP PERHULT JOHNSEN

© ALBIN KARLSSON LANDGREN, 2024.

© PHILIP PERHULT JOHNSEN, 2024.

Supervisor: Daniel Strüber, Computer Science and Engineering

Advisor: Ludwig Friborg, Wiretronic

Examiner: Hans-Martin Heyn, Computer Science and Engineering

Master's Thesis 2024

Department of Computer Science and Engineering

Chalmers University of Technology and University of Gothenburg

SE-412 96 Gothenburg

Telephone +46 31 772 1000

Typeset in L^AT_EX

Gothenburg, Sweden 2024

A Domain-Specific Language for Cross-platform, Edge-deployed Machine Learning Models

A Model Interpretation-based Approach

Albin Karlsson Landgren

Philip Perhult Johnsen

Department of Computer Science and Engineering

Chalmers University of Technology and University of Gothenburg

Abstract

Deploying machine learning (ML) models on edge devices presents unique challenges. The challenges arise from the different environments used for developing ML models and those required for their deployment, leading to a gray area of competence and expertise between ML engineers and application developers. This thesis presents the design and implementation of a domain-specific language aimed at simplifying the deployment of ML models on edge devices, specifically smartphones. It aims to bridge the gap between ML engineers and application engineers, creating a shared platform for deploying ML models on edge devices. The study exists at the intersection of model-driven engineering, machine learning, and cross-platform smartphone development. It explores model-driven engineering in an environment where developers don't have full control over the deployment platform, using model interpretation to generate ML serving pipelines (pre- and postprocessing of data before and after inference) during runtime, thus removing the need to re-release an application upon changes to a pipeline. We follow a design science approach consisting of three research cycles. We elicited requirements through an initial literature study and interviews with engineers at the collaboration company. This was followed by designing and implementing an artifact within the domain presented above. Finally, we evaluated the proposed solution with engineers at the collaboration company through a controlled experiment and subsequent qualitative interviews. The developed artifact consists of a lightweight, JSON-based domain-specific language designed to describe ML serving pipelines, along with an accompanying Flutter library to generate the pipelines during runtime. The evaluation showed that it increased development speed, decreased the amount of code required to make changes to an ML serving pipeline, and made engineers less experienced in mobile development more confident contributing to the domain.

Keywords: Computer, science, computer science, engineering, project, thesis.

Acknowledgements

We wish to thank our supervisor Daniel Strüber for the valuable feedback and guidance throughout the thesis work. We also wish to thank our examiner Hans-Martin Heyn for valuable report feedback during the study.

We are grateful to Wiretronic AB and its engineers for their participation in interviews and experiments, and for providing us with the opportunity to conduct our thesis work at their company and office.

Albin Karlsson Landgren & Philip Perhult Johnsen, Gothenburg, June 2024

Contents

List of Figures	xiii
List of Tables	xv
1 Introduction	1
1.1 Background and Motivation	1
1.2 Study Context	2
1.3 Problem Description	2
1.4 Research Questions	3
1.5 Purpose and Significance of the Study	3
1.6 Thesis Outline	4
2 Theory	5
2.1 Research Gap	5
2.2 Edge-deployed Machine Learning	5
2.3 Cross-Platform Mobile Development	6
2.3.1 ML in Cross-Platform Mobile Environments	6
2.4 Model-Driven Engineering	6
2.4.1 Code Generation and Model Interpretation	7
2.4.2 Metamodels	8
2.4.3 MDE in Edge Devices	8
2.5 Domain-specific Languages	9
2.5.1 Developing a Domain-Specific Language	9
2.5.2 Domain-Specific Languages in the Deployment of ML Models	10
2.5.3 JSON Schemas	10
3 Methods	11
3.1 Design Science Cycles	11
3.2 Cycle 1: Domain Understanding and Initial Artifact Definition	12
3.2.1 Literature Study	12
3.2.2 Repository Analysis and Program Comprehension	12
3.2.3 Interviews	12
3.2.4 Requirements Engineering	13
3.3 Cycle 2: Artifact Design and Development	14
3.3.1 Design and Technology Choices	14
3.4 Cycle 3: Artifact Evaluation	15
3.4.1 Controlled Experiment	15

	3.4.1.1	Metrics	16
	3.4.1.2	Tasks	17
4	Results		19
4.1	Initial Problem Exploration		19
4.1.1	Interview Findings		19
4.1.2	Impact on Artifact Development		20
4.2	Requirements		21
4.2.1	User Stories		21
4.2.2	Functional Requirements		22
	4.2.2.1	Pipeline Specification (DSL)	22
	4.2.2.2	Platform-Specific Model Interpretation (DSL + Architecture)	22
	4.2.2.3	Support Pre-Existing and Custom Operations (DSL)	22
	4.2.2.4	Support Dynamic Changes of the Pipeline (Architecture)	23
4.2.3	Non-Functional Requirements		23
	4.2.3.1	Usability	23
	4.2.3.2	Maintainability	23
	4.2.3.3	Performance	24
	4.2.3.4	Compatibility	24
4.3	Design and Implementation		24
4.3.1	Current Approach		24
4.3.2	Proposed Approach		24
	4.3.2.1	Domain-Specific Language	25
4.3.3	Model Engine		28
4.3.4	DSL Development Tools		29
4.4	Artifact Evaluation		30
4.4.1	Experiment Results		30
	4.4.1.1	Development Time	30
	4.4.1.2	Lines of Code	32
	4.4.1.3	Correctness	32
4.4.2	Hypothesis Testing		33
4.4.3	Requirements		33
	4.4.3.1	Functional Requirements	33
	4.4.3.2	Non-functional Requirements	34
5	Discussion		37
5.1	Research Question 1		37
5.2	Research Question 2		38
5.3	Research Question 3		39
5.4	Cross-Platform Communication		39
5.5	Threats to Validity		40
	5.5.1	Internal Validity	40
	5.5.2	External Validity	40
	5.5.3	Construct Validity	41

6	Conclusion	43
6.1	Limitations	43
6.2	Future Work	44
6.2.1	DSL Expansion	44
6.2.2	Deeper Exploration Within the Domain	44
6.2.3	Expansion to Other Domains	44
 Bibliography		 45
 A Appendix 1 - Initial Interviews		 I
 B Appendix 2 - Experiment Interviews		 III
 C Appendix 3 - Mann-Whitney U Test Code		 V
 D Appendix 4 - Fisher's Exact Test Code		 VII
 E Appendix 5 - Artifact Code		 IX

List of Figures

2.1	An example hierarchy displaying how artifacts, models, and meta-models relate to each other.	8
4.1	The preprocessing method in Java that Wiretronic uses for one of their models.	25
4.2	Abstract syntax of the DSL.	26
4.3	Illustration of how the model engine prepares an ML serving pipeline from a DSL instance.	29
4.4	Comparison of an example JSON schema as defined using Typebox (left) and the actual schema outputted by Typebox (right).	30
4.5	The mean time per task (in minutes) for the old and new approach, respectively.	31

List of Tables

3.1	List of interviewees participating in initial interviews.	13
3.2	Experiment setup.	16
3.3	The engineers from Wiretronic that participated in the experiment, with their respective experience levels.	16
4.1	The time (in minutes) it took for engineers A & B to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.	31
4.2	The time (in minutes) it took for engineers C & D to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.	31
4.3	The lines of code written by engineers A & B to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.	31
4.4	The lines of code written by engineers C & D to complete the subtasks in the first task using the old approach, and the subtasks in the second task using the new approach.	32
4.5	The correctness for engineers A & B when completing the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.	32
4.6	The correctness for engineers C & D when completing the subtasks in the first task using the old approach, and the subtasks in the second task using the new approach.	32
4.7	The startup time (in milliseconds) of the application for the new and old approach respectively.	35

1

Introduction

When deploying an ML model, the choice of target platform can have a significant impact on the ease of deployment. Models are often written, trained, and deployed in a Python environment using libraries such as TensorFlow or PyTorch [1], with the same group of engineers maintaining control during the entire process. However, there are cases where the same toolchain or platform is not available throughout the whole process. One such example arises when deploying machine learning models on edge devices, specifically smartphones. In such scenarios, there is a chance that the training and deployment environments are different and that the deployment environment only contains the already trained models. Thus, the engineers deploying the models to applications might not know how to format inputs and outputs from the model. Therefore, they might land in a gray area of competence and responsibility between those developing and deploying the models. The ML engineers developing the models have most of the knowledge related to the actual models, while the application engineers know the deployment platform, but might lack the crucial context required to effectively deploy the model.

To add to the development and deployment scenario, further complexity arises since the development is happening in a cross-platform environment, a consequence of smartphone applications being developed with two separate platforms in mind (iOS and Android). By bridging the gap between the different engineering roles and simultaneously reducing the amount of equivalent code being written twice, the intention is to improve knowledge-sharing across engineering teams, improve developer experience, and facilitate experimentation and flexibility in development.

1.1 Background and Motivation

The study aims to simplify the process of deploying ML models on edge devices, specifically smartphones. When, for example performing ML inference on image data from a smartphone camera, a series of pre- and post-processing steps is required before and after performing the inference, as different models require inputs of different shapes and output results of different shapes. These steps collectively form a pipeline, referred to as an ML serving pipeline.

Our study and developed artifact aim to facilitate the development and maintenance of ML serving pipelines on smartphones using model-driven engineering (MDE). We propose a domain-specific language (DSL) and accompanying Flutter library that allows developers to easily specify and make changes to ML serving pipelines

deployed on smartphones. The DSL should be able to specify an ML serving pipeline, and the accompanying library should be able to support the execution of the pipeline using platform-specific functionality, based on the contents of a DSL instance. The DSL would act as a platform to facilitate a shared understanding of an ML serving pipeline between the ML engineer and application engineer.

1.2 Study Context

The study is conducted in collaboration with Wiretronic AB and their AI division in Gothenburg, where one of the group members is employed. Wiretronic develops image-based ML products (e.g. object detection and computer vision), including a suite of smartphone applications [2], [3]. As a part of this offering, the company deploys ML models directly on devices, where the variety in architecture and operating systems can make it more difficult to deploy software effectively.

1.3 Problem Description

This problem was introduced by Wiretronic. The applications and corresponding libraries developed at Wiretronic are written in Flutter [4], a library for the Dart programming language enabling cross-platform development of native apps for iOS and Android. Whilst a majority of the functionality can be developed in Flutter, some aspects require writing platform-specific code. This often entails working directly with hardware resources on the device, such as the device’s camera, or hardware-optimized ML libraries [5], [6]. Thus, the developer is required to write equivalent implementations for two platforms, underlining the challenges in deploying ML models on smartphones. Wiretronic believes that since the platform-specific code is constrained to a very specific domain, the workflow can be enhanced. By lifting the development to a suitable level of abstraction and making use of model-driven engineering techniques they could avoid having to write equivalent, domain-specific code for multiple platforms.

The deployment of ML models on smartphones or edge devices, in general, can cause problems with maintenance and updates. When deploying an ML model on a centralized server, the developer can have near full control over that server and perform updates as needed, without users noticing or requiring manual work. Meanwhile, when developing and deploying ML models for smartphones, this process is non-trivial. If the ML model and related functionality are bundled and shipped with the application when installed on a user’s device, we must re-publish the application to the app store of each platform upon making changes or updates and the user must reinstall the application. An alternative method to this is to view the model as an asset to fetch from the application, allowing for easier updates. However, this still requires developers and users to perform the update process if a new version of the ML model requires a different serving pipeline.

In combination, these issues create problems with knowledge-sharing, developer experience, and flexibility when working with ML models on smartphones. It is unclear

who is responsible and most suited to handle the deployment of the ML models, the code often has to be written for two separate platforms, and then subsequently re-deployed for these separate platforms' app stores. By using a DSL and model-driven engineering, we can address communication difficulties, create a single source of truth for the ML serving pipeline, and decrease development as well as deployment efforts. Subsequently, this can also improve the end-user experience since updates to the ML performance can occur without them noticing or having to update the application.

1.4 Research Questions

The study is guided by the findings from studying the role of MDE in edge-deployed ML and applied to the specific context of cross-platform mobile development. Specifically, we will explore how a DSL can be designed and utilized to enhance this development, centered around the following research questions:

- **RQ1:** How can a domain-specific language (DSL) describe an ML model (and, for example, its required inputs, outputs, and pre- and post-processing stages)?
- **RQ2:** How can we best implement and utilize the DSL in a concrete setting, specifically in the development of cross-platform mobile applications?
- **RQ3:** To what extent does the introduction of a DSL and an accompanying library improve the developer experience in the aspects of maintenance, feature development, time-saving, and resource planning?

1.5 Purpose and Significance of the Study

The study aims to answer the questions asked in Section 1.4. The significance of the study is to simplify the process of deploying ML models on edge devices, specifically smartphones. The main challenge identified is the requirement of deploying equivalent solutions to multiple platforms, and we aim to solve this problem using a model-driven engineering approach. The contribution of the study is a DSL for describing the required ML serving pipeline for deploying an ML model on smartphones in a given context. The DSL will support knowledge-sharing across engineering teams, improve the developer experience and, facilitate experimentation and flexibility in development.

While domain-specific languages in machine learning are well-documented, we aim to use this study to contribute to their application in a cross-platform context, which is less explored. Chapter 2 expands on the relevant theoretical background and reviews the state of the art in this field.

1.6 Thesis Outline

Chapter 1: Introduces the thesis project. Explaining our background and motivation, problem description, and research questions regarding this thesis project.

Chapter 2: Consists of the groundwork done to grasp the theory and domain related to the project. Explaining the most significant topics of the research like edge-deployed machine learning, model-driven engineering, and domain-specific languages.

Chapter 3: Explains the design science cycles the project is centered around. How we decided to conduct the work, taking inspiration from design science and requirements engineering defining the requirements and artifact through extensive interviews and casual conversations with the employees at Wiretronic.

Chapter 4: Provides a presentation of our findings during the first research cycle, and how we designed the artifact. How we obtained our insights and how that in collaboration with the background theory resulted in our requirements and implementation.

Chapter 5: Discusses primarily the other possibilities we had in choosing the design of our artifact and our reasoning behind not opting for those paths. Secondly, we discuss potential improvements to our artifact.

Chapter 6: Concludes the thesis, emphasizing a summary of our results and discussing the future work related to our research and artifact.

2

Theory

2.1 Research Gap

We now present several relevant lines of research that, so far, have been developed independently, but their combination has not been considered yet. As of conducting our research, there is currently no available research material about creating a DSL for machine learning models deployed on edge devices in a cross-platform environment. We aim to explore this area, where there are usually several code bases for the ML models that run on Android in addition to iOS.

2.2 Edge-deployed Machine Learning

Edge-deployed ML refers to deploying ML models on edge devices instead of on a centralized server. An edge device can, for example, be a smartphone or Internet of Things (IoT) device, which generally has far simpler hardware than a server in a data center. The deployment of ML models on edge devices has increased significantly in recent years thanks to advancements in both software and hardware [7]–[9]. Despite the less advanced hardware, deploying machine learning models on edge devices presents several advantages when compared to a centralized approach. Transmitting potentially sensitive or private data to a remote server introduces the risk of data leakage, with a fault in the remote system potentially leading to personal or financial consequences [8]. Additionally, eliminating the need for connecting to an external service for ML inference can improve both latency and reliability, as the potential bottleneck introduced by a weak network connection is removed. Despite these improvements, deploying ML models on edge devices, specifically smartphones, is not straightforward. One reason for this is the heterogeneity of underlying architecture [10]. A wide selection of libraries to deploy ML models on smartphones exists, but they each perform differently depending on the device’s hardware configuration. A difference in cache size or GPU capacity can cause two libraries accomplishing equivalent tasks to perform differently, and with the wide range of hardware configurations present in the market, it is difficult to develop a solution optimal for every device [10]. Additionally, opting to deploy an ML model on devices instead of in a centralized environment can create obstacles to improvement and maintenance. When deploying an ML model in a centralized environment, the developer has full control over software and hardware and can develop the artifacts surrounding the model for a single environment. If the development is instead targeted at smart-

phones, the model has to be deployable on both iOS and Android devices, which each have distinct underlying architectures for deploying custom ML models [5], [6].

2.3 Cross-Platform Mobile Development

Mobile developers targeting both iOS and Android users may opt for a cross-platform framework, which enables the creation of separate, native builds for both platforms from a single codebase. The two most widely used frameworks for this purpose are Flutter and React Native [11]. Flutter is an open-source framework for the Dart language maintained by Google, while React Native is a JavaScript framework maintained by Meta. While there are some differences in architecture, both frameworks abstract away platform-specific details, allowing developers to focus on a single, platform-agnostic codebase. This abstraction layer translates the shared code into platform-specific components, aiming to achieve native functionality and performance for both iOS and Android. When requiring access to platform-specific features, often hardware, developers can opt to write native code for each platform. Although React Native has a new architecture in development that will allow for easier communication between the cross-platform and native layers, current implementations of both Flutter and React Native require serialization of data for inter-layer communication [12], [13].

2.3.1 ML in Cross-Platform Mobile Environments

In the deployment of machine learning (ML) models within cross-platform mobile environments, it is often advantageous to utilize platform-optimized ML frameworks. An example of such a framework is Core ML for iOS [5]. The advantages presented by such an approach increase as the ML model requires interaction with other hardware functionalities, such as the camera of the device. In a scenario where continuous inference on a camera stream is required, significant processing time can be saved by performing the entire computation flow on the native layer, as it omits the data serialization introduced by inter-layer communication [12].

2.4 Model-Driven Engineering

Models are important in many scientific contexts to understand the basics of a field or domain. They are descriptive of a system or prescriptive for determining the scope or details of a problem. This is no different in software development, adopted in model-driven engineering (MDE). MDE reshapes software engineering by emphasizing high-level abstraction and model-centric approaches. The abstractions offered by MDE facilitate easier adaptation of new technologies. This abstracts away platform-specific details, making it viable when working in cross-platform environments [14].

MDE employs model-based approaches that can improve the daily practice of software professionals [14]. Through the focus on creating and examining models, these aspects can capture various aspects of a problem-specific domain [15]. This allows

for complex systems to be more understandable and more easily translatable across different platforms through code generation or model interpretation [16].

Adjacent to MDE are concepts such as Model-Driven Development (MDD) and Model-Driven Architecture (MDA). MDD focuses on placing models at the center of the development process, and implementations of these models often are either fully or partially generated. MDA is a term first coined and used by the Object Management Group (OMG) to narrow the focus of their OMG standards to modeling and transformations. MDD serves as a superset to MDA, and MDE in turn acts as a superset that encompasses both MDD and MDA. MDE offers a broader perspective and application of model-based methodologies in software engineering. In addition, MDE is sometimes referred to as Model-Driven Software Engineering (MDSE), but this term is merely a synonym with MDE, both encapsulating performing software development through abstractions and modeling [14].

2.4.1 Code Generation and Model Interpretation

There are several approaches within MDE to go from a model to executable software. One such approach commonly applied is code generation, where a model is transformed into a program in a suitable language that can subsequently be executed. This process allows for developer intervention where needed since the generated code can be edited upon generation. Furthermore, as a consequence of the code being generated before execution, this approach does not introduce any run-time overhead. However, while avoiding performance overhead, a disadvantage of code generation is having to re-generate and re-deploy the software if making changes to a model [14].

In addition to code generation, another use case for automating software development is model interpretation [14]. Model interpretation does not generate code, it instead implements a generic engine, e.g. a library, that parses and executes the model on the fly. This comes with several advantages as noted by Brambilla *et al.* [14]: it allows making changes to the model or engine without an added code generation step, easier portability between platforms, and not having to interact directly with the source code.

There are also a few concerns about model interpretation. As it is a black-box approach, the application will be dependent on the library or tool serving as the engine. If this implementation is not written optimally it can cause performance issues that are difficult to solve, as the developer may not have insight into how the engine operates. While this is a common reason for discarding this approach, it is not a concern for most applications [14].

While model interpretation is often viewed as an alternative to code generation within MDE, they are not exclusive alternatives. The two techniques can often be used in a hybrid manner [14]. For example, developers may use model interpretation during development for faster prototyping, while using code generation for production to eliminate runtime overhead or minimize bundle size.

2.4.2 Metamodels

Previously, we have described a model within MDE as an abstraction of how some software artifact is constructed. If moving one level of abstraction higher, we can define a *metamodel* as a high-level abstraction of a model [14], [15]. Designing a metamodel helps guide developers to have a structured approach when designing a model's intended properties and behavior in a software system.

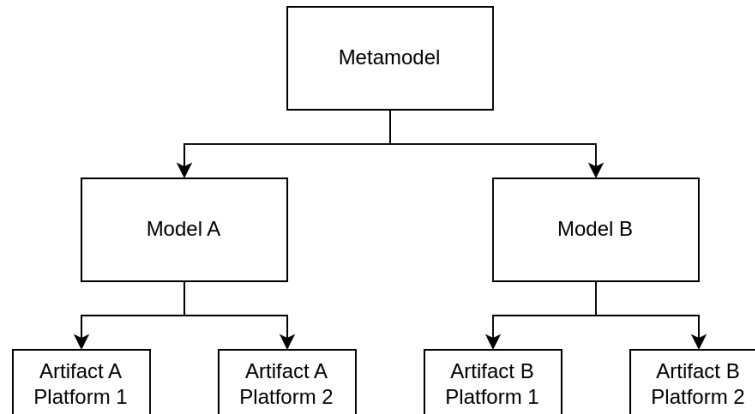


Figure 2.1: An example hierarchy displaying how artifacts, models, and metamodels relate to each other.

Figure 2.1 displays how artifacts conform to a model and how this model can conform to a metamodel. Instead of focusing on a specific platform, the model helps describe the respective artifacts on a higher level, and these models are in turn guided in their shape by the metamodel, describing how all models should be structured.

2.4.3 MDE in Edge Devices

As hardware improves and new areas of applicability arise, the demand to deploy ML models on edge devices increases. However, integrating ML models into edge device environments still comes with many limitations in terms of computational resources, power constraints, and network communication [17].

Furthermore, there is a significant heterogeneity in edge devices, spanning from low-memory microcontrollers to high-end smartphones. Working in this domain can therefore require familiarity with several techniques and operating systems.

Within the specific domain of mobile development, Vaupel *et al.* [18] discuss how model-driven techniques can be used to create flexible, cross-platform mobile applications, stating that models should be "As abstract as possible and as concrete as needed." [18], [19]. By opting for model-driven techniques and using higher abstraction levels we can create separate native builds from a single source, similarly to techniques mentioned in Section 2.3.

2.5 Domain-specific Languages

DSLs are languages used within a particular domain or problem, unlike general-purpose languages like Python, C++, or Java. These DSLs are designed with domain-specific abstractions and notations, making them more concise and accessible, and thereby easier for providing solutions within the domain [20]. MDE plays a significant role in the development of such domain-specific languages, as they often differ from general-purpose languages. In many cases, MDE is employed specifically for working with DSLs [21]. When defining a DSL, the terms abstract syntax and concrete syntax are often used, and they are relevant in this study. The abstract syntax refers to how the DSL is modeled, i.e. which features are available and how they combine and interact with each other, while the concrete syntax refers to the actual grammar of the DSL and how the source files are written [14]. While the concrete syntax is represented through the DSL itself, the abstract syntax is often represented through metamodels. As Paige *et al.* [22] states, a metamodel can be defined as the "description of the abstract syntax of a language, capturing its concepts and relationships, using modeling infrastructure".

2.5.1 Developing a Domain-Specific Language

The development of a DSL requires a structured approach [23]. Firstly, the necessity of a DSL has to be established, followed by an analysis of the domain to elicit the specific requirements, and thereafter, designing the DSL syntax and semantics to meet the requirements. Implementation of the DSL requires appropriate tools and technologies to leverage its advantages. Lastly, testing and validation to ensure the DSL efficiently solves the domain-specific challenges it was developed to solve.

Developing DSLs can be done in both Language Workbenches (LW), like Xtext [24], GEMOC, and MetaEdit+, or by using more lightweight approaches, such as a JSON Schema [25]. In the literature review done by Korani *et al.* [26] from 2023, Xtext is still the most used language framework for developing textual languages, this is also backed by Moin *et al.* [27] that uses the Xtext framework for developing ML-Quadrat [28], an open-source prototype based on MDE for full code generation on IoT devices. Using LWs enables more comprehensive tooling and support for editors to make it work like a programming language. However, the LWs require more time to develop and developers have to learn a new programming language. On the other hand, using a JSON Schema enables defining specific data structures and can be written in JSON or YAML. Modern IDEs support features such as autocomplete and type-checking in even lightweight tools.

Fundamentally, there are two ways for DSLs and regular code to interact, either as external or internal DSLs. External DSLs keep the DSL and GPL code in different files, where the DSL is transformed into a programming language and executed, like in Xtext. SQL is a well-known external query DSL used in relational databases [29]. Internal DSLs have both types of code in the same file, reusing the same grammar as the General Programming Language (GPL) it is written in cohesion with.

2.5.2 Domain-Specific Languages in the Deployment of ML Models

DSLs can play a significant role in the deployment of ML models, especially on edge devices where computation power and memory are limited. A DSL can help ensure type and function compatibility, which is an integral part for models used for tasks such as image recognition and text processing. In addition, providing the ability to efficiently manage tasks such as inputs and outputs. The paper by Zhao *et al.* [30] introduces a system that exemplifies the use of a DSL in such a context.

Beyond the development and deployment, MDE has also been used to support the management of orthogonal ML aspects, such as asset management [31] and dataset management [32]. Traditional version control systems (VCS) can struggle to handle complex assets such as ML models and datasets. In the paper by Idowu *et al.* [31], they address these asset management challenges by introducing the Experiment Management Meta-Model (EMMM). A meta-model to characterize ML asset structures as concepts and their relationships observed in state-of-the-art tools, and conceptual VCS structures that can hold both ML and traditional assets. Meanwhile, Giner-Miguel *et al.* [32] presents DescribeML, a tool for utilizing a DSL to describe datasets. This tool aims to enable a more data-centric approach in ML, to handle issues like undesired model behaviors resulting from biased predictions.

2.5.3 JSON Schemas

JavaScript Object Notation (JSON) can be used to define DSLs, as DSLs can be embodied as configuration files in applications [33]. JSON is a data serialization format widely adopted to either store data physically or transfer it over the internet [34]. It is a semi-structured document format, that is possibly the most popular format for data exchange over the internet [35], [36]. It allows developers and IT professionals to transfer data structures across programming languages and environments, without having to worry about said environments. Instead, the data can be serialized or parsed in any language.

JSON Schemas and JSON documents differ in their purpose. A JSON document contains the data to be sent or stored, organized in JSON objects. Meanwhile, JSON Schemas are used to define the structure of which a JSON document should adhere to, to ensure compatibility and consistency. The schema can then also be used to validate a JSON document. [37]

JSON Schemas are the standard schema language for structuring JSON data. It is based on a combination of structural operators that describe values, arrays, and objects, with logical operators like negation, conjunction, and disjunction [38]. JSON Schema validators have been developed for many programming languages and they are used to make software and data transfer more reliable [34].

3

Methods

This study employs a design science approach, combining theoretical research with a practical application of the research into a specific artifact. Firstly, we conducted research into the potential role of MDE in the deployment of ML models on edge devices. We applied these findings to develop an artifact to solve the specific problem of deploying ML models in cross-platform mobile environments. The problem is being addressed through the three aforementioned research questions and three distinct cycles explained in Section 3.1.

3.1 Design Science Cycles

Following the process presented by Knauss [39], each cycle was centered around a specific phase in the design science process while iteratively advancing the understanding and progress of each research question. Below is a summary of the work conducted in each cycle of the study.

- **Cycle 1 (RQ1):** This cycle was about a major focus on research and understanding of the domain. We studied existing literature covering topics relevant to the study. Developed small proof-of-concept solutions and evaluated their viability in this context. We identified clear requirements for our artifact, both functional and non-functional, to deepen our knowledge about the domain and Wiretronic’s needs. We employed tools such as interviews, frequent conversations with employees, and inspection of source code.
- **Cycle 2 (RQ2):** This cycle primarily focused on applying our findings through the concrete implementation of the artifact, iteratively verifying the development against the requirements specified by Wiretronic. Complemented with research into our domain and specific implementation details.
- **Cycle 3 (RQ3):** We evaluated our developed artifact by conducting a two-fold evaluation. Firstly, we conducted an internal evaluation, focusing on the requirements defined in cycle 1. Secondly, an external evaluation in collaboration with our supervisor from Chalmers and the respondents at Wiretronic. To test the suitability of our artifact we conducted a controlled experiment at Wiretronic, where two groups performed a set of tasks using the existing approach and the new approach. The two groups were measured with respect to time, lines of code, and correctness. Additionally, a final interview to evaluate their experiences of implementing cross-platform specific code using our new

approach compared to the old approach.

3.2 Cycle 1: Domain Understanding and Initial Artifact Definition

As stated, we spent most of the first cycle researching the domain and its specific representation at Wiretronic, using this information to help define our requirements and the scope of the study. This section covers the methods applied during this phase.

3.2.1 Literature Study

The literature study was conducted as a foundation for our interviews and subsequent requirements elicitation. During cycle 1, the literature review provided insights into various design options for our artifact. We explored developing the DSL using Xtext as a language workbench or the more lightweight approach of JSON or YAML. We assessed whether Wiretronic would benefit more from code generation or model interpretation, and connected to this, we explored the benefits and drawbacks of a build time respectively runtime approach to our model-to-code transformation, considering feasibility within our time frame. From our interviews in Section 3.2.3 we got clarifying answers that helped us decide our final approach.

3.2.2 Repository Analysis and Program Comprehension

We spent part of the first cycle examining an existing library developed at Wiretronic. This library powers all of Wiretronic’s machine learning operations on edge devices, here limited to smartphones. This was primarily done as a program comprehension activity, as we required a thorough understanding of the domain and current state of development to have informed discussions with engineers at the company, identify constraints for future requirements elicitation, and find potential areas of enhancement.

The applications are written in Flutter, using Java for Android-specific functionality and Swift for iOS-specific functionality. This resulted in large parts of the program comprehension being conducted twice, as the machine learning code was implemented both in Java and Swift. This gave us two possibilities to understand most of the relevant code instead of one. Aside from serving as a tool to inform our requirements elicitation and development, analysis of the library was used as a tool to better understand the Flutter architecture and how it handles communication between the cross-platform and native layers.

3.2.3 Interviews

At Wiretronic, there were two engineers with experience in applications and libraries relevant to the study, therefore these two were chosen for qualitative interviews. The interviewees had experience both in developing the ML models and deploying them

on devices, which allowed us to obtain a holistic view of their current processes with a small number of interviewees. The interviews, performed as part of our problem exploration in the first cycle, aimed to document workflows and challenges to inform our subsequent requirements elicitation. While the interviews were conducted during the initial phase of the study, they were not conducted immediately. We deemed it necessary to first grasp the theoretical concepts relevant to the study, in addition to performing program comprehension. This was done to ensure we would go into the interviews well-informed and that they would serve their purpose.

Label	Role	Platforms	Experience
Interviewee A	Engineer/System Architect	ML + iOS	4 years
Interviewee B	Engineer/System Architect	ML + Android	4 years

Table 3.1: List of interviewees participating in initial interviews.

The interviews were conducted in a semi-structured format to obtain both quantitative data, such as tools currently in use, and deeper, qualitative insights into workflows and challenges. To allow for this structure, we used a set of pre-defined questions available in Appendix A in combination with follow-up questions to elicit more detailed information. Each session began with a standardized introduction to maintain consistency across interviews, regardless of the participants' prior knowledge of the study. When crafting the interview questions, we deliberately included some questions where we expected to already know the answer. This measure was taken to ensure that basic needs, or *must-be* requirements, were not overlooked. These requirements are often taken for granted and go unnoticed if fulfilled, but failing to fulfill them can render the artifact unusable [40].

Upon conducting the interviews, we analyzed them as part of our requirements identification. Since the majority of the interview content was equivalent across the two interviews, they could be directly compared to identify challenges and pain points identified by both engineers. We also followed up the interviews with informal discussions, helping us draw conclusions and inform requirements when opinions or statements presented in the interviews were conflicting.

3.2.4 Requirements Engineering

We used the insights obtained from examining relevant repositories at Wiretronic and interviewing engineers when specifying our requirements. Analyzing repositories gave us a good overview of their existing systems and possible areas of improvement within our scope. Additionally, the interviews were valuable in providing context to our findings, ensuring that they align with the requirements and priorities of the company. The requirements identification was initiated by developing a set of user stories centered around a persona representing a developer at Wiretronic. This helped us bridge the requirements elicitation and requirements identification phases, consolidating the information we had obtained without getting caught up in implementation details.

After defining a set of suitable user stories we began defining requirements, both

functional and non-functional, rooted in the user stories. Naturally, this phase helped in setting up a set of more concrete, measurable goals for the project. Defining the requirements was an important tool in defining the scope of our study and creating a mutual understanding of priorities among ourselves and with the engineers at Wiretronic. Furthermore, the requirements were vital for the third and final cycle of the study, when performing validation and verification of the developed artifact.

3.3 Cycle 2: Artifact Design and Development

The purpose of the second cycle was two-fold: it first involved transforming the data collected in the first cycle to well-informed technology and design choices, and secondly, it involved designing and implementing the artifact. This section covers how we conducted this transformation, as well as how the design and implementation phase was conducted.

3.3.1 Design and Technology Choices

This section and the choices we made can be divided into two separate areas making up our artifact: the DSL and the accompanying library.

When designing the DSL, the primary guiding factor was the interviews with engineers, as no similar project had been conducted at Wiretronic before. The interviews, informed by the initial research activities, helped us narrow down what specific purpose we should aim to solve. The technologies chosen for the library are primarily rooted in practices already in place at Wiretronic to avoid obstacles in the handover of the artifact at the end of the study and to ensure compatibility with relevant applications. This information was elicited through the interviews and our analysis of existing repositories.

The possibility of using a JSON Schema to define the DSL was explored before the actual study began. Through the interviews and subsequent requirements engineering it was deemed a viable and preferable option during the first research cycle. We found that the main role of the DSL would be to describe an ML serving pipeline, and not write the actual implementation and logic, thus making a JSON Schema a fitting choice. After this decision was made, more focus was put into how to best describe the model metadata and the pre- and postprocessing steps. This involved going through the existing models and comparing which aspects of the current ML serving pipelines that are shared, and which are unique for one or a set of specific models.

After identifying the required content of the DSL, the concrete syntax had to be established. Thanks to the lightweight nature of JSON Schemas in contrast to developing a programming language, we were able to iterate on the syntax quickly and try out several variations of the syntax in a single working session. Additionally, some of the choices of this process were made automatically due to limitations of the JSON specification, as highlighted in Section 4.3.2.

The library was designed in parallel with the DSL, ensuring both that any additions

or changes made to the DSL would be feasible to implement in the library and that we could find a suitable place for them. When designing the functionality for preparing and running the actual ML serving pipeline, we made several choices based on our initial research and the interview feedback.

It was clear that, since a camera-based ML application can receive data as a stream of images, the overhead introduced by our library must be minimal. This meant that we wanted to avoid parsing the DSL instance each time, and also avoid conditional statements during execution, based on the parsed DSL instance. Thus, we implemented the pre- and postprocessing as a series of individually contained *steps*, all implementing an interface with the necessary method stubs. Thus, the pipeline lists consist of generic pre- and postprocessors, and not the concrete implementation, according to the dependency inversion principle [41].

This helped us to separate the preparation and execution of the pipeline, as all steps had a method for setting it up with all the correct parameters and a separate method for executing it. While this was mainly done to eliminate any DSL-related logic during execution, it also helped when designing the functionality to implement custom pre- or postprocessing steps. By allowing the developer to implement an anonymous class implementing the interface directly in the consuming application, they can be confident that the step will be compatible with the pipeline, as long as the implementation is fault-free. In MDE, this functionality is considered part of the model engine, which is presented in more detail in Section 4.3.3

3.4 Cycle 3: Artifact Evaluation

This section covers our evaluation of the developed artifact, which was the main focus of the third and final cycle. Here, we first conducted an internal evaluation of the developed artifact, comparing the result with the visions presented by Wiretronic and the set of requirements we developed as a result of our initial exploration. Secondly, we evaluated the artifact together with Wiretronic, performing a controlled experiment with two groups of engineers. In doing this evaluation, we covered both aspects of verification and validation, ensuring not only that the artifact has been built correctly, but also that it solves the correct problem.

3.4.1 Controlled Experiment

To perform the evaluation, we performed a controlled experiment. The purpose behind this was two-fold: first, we aimed to identify the specific impact of our artifact, and second, maintaining increased control over the experiment helped ensure similar conditions for each trial, minimizing the impact of outside factors. The experiment was carried out using a Latin square design [42], where two groups each performed two tasks. One group used our artifact to solve the first task and not for the second task, and vice versa, as is displayed in Table 3.2. All sessions were performed in a 60-minute time slot, ensuring all participants had the same time to perform the tasks. Furthermore, the two groups received identical presentations and documentation for our artifact. Because of the small available sample size, we

utilized stratified sampling [43]. The engineers were categorized into two groups: experienced and inexperienced, with both experienced engineers having four years of experience and inexperienced zero, not having worked in the environment at all. We then formed the two experiment groups with an equal number of experienced and inexperienced engineers, as visible in Table 3.3. The Latin square design aimed to minimize the learning bias that comes from the participants improving their performance by repeating similar tasks. By alternating the order in which the tasks are performed, and the use of the artifact across both groups, we can effectively reduce the bias. Group 1 will be the control group for task 1, and the treatment group for task 2, and vice-versa for group 2.

	Group 1	Group 2
Task 1	not using artifact	using artifact
Task 2	using artifact	not using artifact

Table 3.2: Experiment setup.

Group 1	Engineer A**	Experienced
	Engineer B	Inexperienced
Group 2	Engineer C*	Experienced
	Engineer D	Inexperienced

Table 3.3: The engineers from Wiretronic that participated in the experiment, with their respective experience levels. *Interviewee A, **Interviewee B.

With this experiment, we aimed to identify whether the introduction of our artifact improves the workflow of the specific process it is designed to improve, to give answers to **RQ3**. Opting for a contrived setting allowed us to identify the impact of the artifact, albeit at the cost of generalizability and realism [44]. To provide further nuance and compensate for the drawbacks of a controlled experiment, we conducted semi-constructed interviews with the participants to gain qualitative insights.

3.4.1.1 Metrics

In RQ3 we wanted to answer to what extent our artifact improves the developer experience in aspects such as maintenance, feature development, time-saving, and resource planning. We used the experiment to obtain quantitative data and combined this with the interviews for qualitative data. The quantifiable metrics we observed through the experiment were the following:

- Time per completed task. Measured in minutes, extracted from commit timestamps.
- Lines of code written to complete each task. Measured in lines inserted and lines deleted for each commit.
- Correctness, a binary metric of whether the task was performed correctly or incorrectly. Measured by manual static analysis of the solution, and occurrence of runtime errors after the experiment.

The post-experiment interviews helped us obtain qualitative data about more subjective metrics, identifying how usable, intuitive, and useful the artifact can be for the engineers' daily workflow. The questions asked in these interviews are available in Appendix B.

Additionally, we performed hypothesis testing on our metrics, specifically time and correctness to get a more comprehensive view of our results. We expected the data to not be normally distributed due to a small sample size, natural variations in human performance, and the variability in the experience levels of the engineers. For the development time, we utilized the Mann-Whitney U Test [45]. This test is suitable because it is non-parametric and does not assume a normal distribution, making it appropriate for small sample sizes and discrete time data. For the correctness metric, we used Fisher's Exact Test [46]. It is designed for categorical data, in our case correct and incorrect, and is ideal for small sample sizes.

The hypotheses for the Mann-Whitney U test:

- Null Hypothesis (H0): There is no statistically significant difference in the development time between the old and new approaches.
- Alternative Hypothesis (H1): There is a statistically significant difference in the development time between the old and new approaches.

The hypotheses for Fisher's Exact Test:

- Null Hypothesis (H0): There is no statistically significant difference in correctness between the old and new approaches.
- Alternative Hypothesis (H1): There is a statistically significant difference in correctness between the old and new approaches.

3.4.1.2 Tasks

As stated, we designed two example tasks to evaluate the artifact. Task 1 had three subtasks and Task 2 had two subtasks. These were designed with the pain points of Wiretronic in mind, identifying how effective the artifact can be in maintenance for both the pre- and postprocessing parts of an ML serving pipeline. Therefore, Task 1 is completely related to preprocessing and Task 2 is completely related to postprocessing.

Task 1 - Assessing preprocessing: Given an existing model with accompanying pre- and postprocessing methods implemented, the engineers will perform the following subtasks:

- Change the path from which the model is loaded.
- Modify the size of the input data, that the image will be resized to from 300 by 300 to a new specified dimension, 380 by 380.
- Enable normalization for the input image.

Task 2 - Assessing postprocessing: The model that has the least trivial post-processing is a multi-headed model used for several computer vision tasks. Being multi-headed, it can both provide e.g. whether an item is visible in the frame, and

3. Methods

produce a bounding box for locating the item.

- Adjust the threshold of the binary classification head named *is_visible* to 0.5.
- Implement interpolation for the binary classification called *size*. Set the size to 300 if below the threshold, otherwise set it to 500.

4

Results

This chapter presents the findings of our study, the subsequent artifact implementation, and the evaluation of the artifact. It lays out the requirements that guided the artifact implementation and evaluation, along with the reasoning behind each requirement.

4.1 Initial Problem Exploration

This section is dedicated to presenting our findings from the first cycle, focused on defining the artifact. This entails our literature study, repository analysis, and interviews. The literature study primarily focused on **RQ1**, identifying how we can implement the DSL for this specific scenario. Meanwhile, the repository analysis and interviews were aimed at exploring **RQ2** and **RQ3**, identifying how the introduction of a DSL for ML pipelines can improve development processes within Wiretronic.

4.1.1 Interview Findings

We primarily obtained insights into existing development processes and potential enhancements through the engineer interviews. Interviewee B stated that a DSL and accompanying tools would help in the development and testing of ML serving pipelines, specifically for iOS. Stating that since he does not use MacOS, a requirement for building iOS applications in Swift, he can not currently develop for iOS. Instead, if making changes to an ML serving pipeline, he would have to write and test the changes in Java and then pass development to Interviewee A, who can implement the equivalent functionality for iOS in Swift. He mentioned that with a DSL he could instead define an ML serving pipeline using the DSL and then be confident that the iOS implementation will work, as long as the DSL instance is written correctly. Interviewee A independently pointed this out as well, underlining the fact that native development and related communication are obstacles in their current workflow. Furthermore, the two engineers agreed that an additional problem they would like to solve is having to publish a new version of the library when either making a change to an ML serving pipeline or implementing a new model.

When asked about the language design, Interviewee B stated that they would prefer writing the pipeline steps in a format they are familiar with and can get used to quickly rather than a completely custom DSL since there are only two platforms. They used the reasoning that if they were to learn a new language or platform, they

could learn the other platform (in their case, iOS/Swift) instead of a new DSL.

The two interviewees presented slightly different approaches to implementing the DSL in an application. Interviewee B suggested that it could be a part of the build process, i.e. generating platform-specific code for the ML pipeline when compiling the application. Interviewee A, however, noted that he would prefer that the DSL be bundled with the application, loaded and parsed during runtime, and then used to configure the pipeline. This suggestion can be classified as a model interpretation-based approach, as it parses and executes a model during runtime [14]. This process requires including all possible pipeline operations in the application bundle. He stated that the performance implications would be negligible, especially in comparison to loading an ML model from either the disk or over the network, which is already done in the applications. The suggestion by Interviewee B would satisfy their shared pain point of having to republish the library when making a change to an ML serving pipeline, but it would still require publishing a new version of the application. Interviewee A’s suggestion would also remove this step, but it could prove less flexible if a developer needs to add currently non-existent functionality or functionality not general enough to be part of the library.

4.1.2 Impact on Artifact Development

Here we present some decisions made after conducting our initial studies and interviews. While Section 4.3 explores the design and implementation of our artifact in more detail, this section aims to provide relevant context for Section 4.2, which lays out the requirements guiding the artifact development.

When re-examining the problem after our research and interview study, we decided to opt for an approach based on model interpretation. This decision was primarily driven by two factors. Firstly, the interviews along with further discussions with engineers confirmed that the set of operations used for image transformation is limited and overlaps significantly across pipelines, confirming our previous findings from examining repositories. This highlighted that the configuration of arguments would benefit more from abstraction than the development of completely new functionality. Secondly, by opting for a model interpretation-based approach, the need to release a new library or application version upon making changes to the pipeline is removed, as previously highlighted. Instead, the pipeline can be updated dynamically, for example by fetching it from a remote server, thanks to the required functionality being bundled in configurable modules with the application.

While it seems suitable for this scenario, opting for a model interpretation-based approach may bring drawbacks. As highlighted previously, if a new ML model that requires custom preprocessing functions is introduced, this functionality will not be present in the library. In this case, the DSL and library either have to be extended to include this functionality, or we would need to include a way for a developer to reference one-off functions residing in the application in the DSL with suitable syntax. This does in turn introduce a problem of runtime safety. If we fetch a new DSL instance and this includes functionality not present in the application, the pipeline will not be configured correctly.

Upon discussions with engineers, we still deemed the model interpretation-based approach to be most suitable. If using code generation and avoiding runtime configuration, completely new functionality not supported by the DSL would still require substantial maintenance work and manual updating of either the library or applications consuming it.

As highlighted by our interview study, the DSL needs to be easy to learn and adopt compared to mastering a new platform. This, in combination with our specific context of defining a machine learning serving pipeline using pre-defined functionalities, we decided that employing a JSON Schema for the DSL would be an effective approach. This choice seemed more suitable than opting for a more complex and advanced tool like Xtext since the primary goal is to describe pipeline steps and we do not require more detailed application logic within the DSL. Utilizing a JSON Schema offers several advantages: it simplifies the versioning of the DSL and allows for the validation of DSL instances against the schema. These validation abilities in turn provide syntax highlighting and integrated documentation within the developers' editors for increased usability and ease of adoption.

4.2 Requirements

This section will present the requirements identified through our requirements engineering process, presented in further detail in Section 3.2.4. This entails both the user stories, focused on creating a high-level view of the solutions provided by our artifact, along with our functional and non-functional requirements. The requirements are presented together with a short description aimed to provide further context and reasoning behind the requirement.

4.2.1 User Stories

User stories are features written from the perspective of a user, in our case a developer [47].

- **UC1:** As a developer, I want to be able to create and modify ML pipelines for multiple platforms without requiring platform-specific knowledge.
- **UC2:** As a developer, I seek to avoid writing equivalent, platform-specific code for multiple platforms when deploying ML models.
- **UC3:** As a developer, I want a configuration file in a format I recognize, like JSON, to quickly change ML model parameters for rapid experimentation to enhance efficiency.
- **UC4:** As a developer, I aim to dynamically adjust ML model configurations using the DSL at runtime, thus avoiding releasing new application or library versions for changes to the configuration.
- **UC5:** As a developer, I wish to use pre-built templates for common ML tasks, enabling me to concentrate on developing new and unique features for

improving model performance.

- **UC6:** As a developer, I need a framework to easier identify potential failures in the ML pipeline, reducing manual debugging efforts.

4.2.2 Functional Requirements

The functional requirements specify the functions of the system, the features it is going to have, and how it handles data [48].

4.2.2.1 Pipeline Specification (DSL)

The ML serving pipeline refers to the set of processing steps required for an ML serving model. Each step in the process is a pipeline step, that performs a specific operation or transformation to data. As specified in **FR1.1**, this would include the pre- and postprocessing steps required before and after using the ML models. The preprocessing steps Wiretronic uses include cropping an image, rotating an image, changing image format, normalizing pixels, and initializing buffers for storing image data. The post-processing steps include tensor conversion, and extracting tensor data into other formats.

- **FR1.1:** The DSL should be able to specify which pre- and postprocessing steps are required for an ML model in a given context.
- **FR1.2:** The DSL should be able to be validated against a JSON Schema to ensure its correctness.

Given the need for a clear and flexible way to define these pipelines, we have chosen to use JSON Schemas for our DSL. JSON Schemas provides a structured yet lightweight approach to defining the syntax and validation rules for our DSL, ensuring compatibility and ease of use across different platforms.

4.2.2.2 Platform-Specific Model Interpretation (DSL + Architecture)

When specifying the steps in the DSL, the library should allow for model interpretation directly in Swift and Java. It ensures the application can be run across different platforms, in this case iOS and Android, by abstracting away the complexities of writing platform-specific code, while also allowing for changes to the ML serving pipeline on the fly.

- **FR2.1:** The DSL should enable model interpretation in Swift and Java, initiating an ML serving pipeline from existing native functionality based on the steps defined in an instance of the DSL.

4.2.2.3 Support Pre-Existing and Custom Operations (DSL)

A tool like this needs to be able to maintain the freedom of implementing specific operations if needed. Our tool already provides the existing operations mentioned in Section 4.2.2.1, however, these are still pre-defined operations Wiretronic uses for their ML models. When working with ML models the preprocessing steps can

significantly impact the predictions of the models, hence making it an iterative process using different operations that could need these custom operations [49].

- **FR3.1:** The DSL should enable the developers to use local functions instead of those pre-defined in the DSL.

4.2.2.4 Support Dynamic Changes of the Pipeline (Architecture)

One of the advantages of implementing a DSL and library solution is that it enables dynamic changes during runtime. By having the ML serving pipeline set up dynamically through a configuration JSON file, we can change the model serving parameters without Wiretronic having to release new versions of their library. Since all functionality already exists in the library, we can dynamically load new model parameters when changes happen to the configuration file, or initialize a new configuration file.

- **FR4.1:** Being able to switch between several configurations while the application is running, enabling A/B testing of pipelines.

4.2.3 Non-Functional Requirements

Non-functional requirements, or quality requirements, specify how well the system performs its functions. It is very important to address these alongside the functional requirements, as they play a crucial role in what we want to achieve with the requirements as stated in Section 4.2 [48].

4.2.3.1 Usability

Usability refers to how friendly the system is to users [50]. The artifact aims to ease the workflow of the developers, hence it needs to be intuitive and have a low learning curve.

- **NFR1.1:** The system should be easy to learn, allowing developers to use it with minimal training required.

4.2.3.2 Maintainability

Maintainability here refers to the ability to improve and understand software [50]. As the thesis project is in collaboration with Wiretronic, it is essential to make the artifact easy to build further upon by the company after the completion of the project. By writing documentation about our solution, the developers at Wiretronic should easily be able to understand our library and DSL to make changes or add new features.

- **NFR2.1:** The system should be easy to update, with clear documentation and guides.
- **NFR2.2:** It should facilitate the addition of new ML serving pipeline features without having to make substantial modifications to the existing code.

4.2.3.3 Performance

Performance defines how fast a software system or component responds to actions [50]. In Section 2.4.1 it says performance may be a concern for some when using model interpretation. Through our research and implementation, we aim to prove that using model interpretation should not negatively impact the application startup time when initiating the ML serving pipeline through the library and DSL.

- **NFR3.1:** The system should not add more than 50ms to the application startup time when initiating an ML serving pipeline from an instance of the DSL.
- **NFR3.2:** The system should not cause performance overheads when running an application containing an ML serving pipeline dynamically set up by the library.

4.2.3.4 Compatibility

Compatibility refers to a system that exists and interacts with another system in the same environment [50]. As the system is in a cross-platform environment it is important to not have any limitations due to different operating systems or IDEs.

- **NFR4.1:** The system should work across multiple platforms (MacOS, Windows, and Linux).
- **NFR4.2:** The system should work in Flutter codebases.

4.3 Design and Implementation

4.3.1 Current Approach

Figure 4.1 is a code snippet from the existing library at Wiretronic, it displays how the preprocessing is written in Java for one of their models. The method performs cropping, rotation, and normalization of an image, with the parameters for image size being instance variables in the Java class. When implementing a new ML model or making changes to an existing ML serving pipeline, the developers will also have to write this code for Swift to support iOS devices. As will be presented in this section, our DSL abstracts away the platform-specific details and provides the developer with a single interface to specify the ML serving pipeline.

4.3.2 Proposed Approach

In this section, we propose an alternative approach to manage the ML serving pipelines in cross-platform mobile environments, decoupling this configuration from the underlying platform. This proposal is the result of the previously outlined requirements definition and the work done to inform that. It consists of two separate but connected parts: the DSL aimed to aid developers in specifying the ML serving pipelines in a single, familiar format, and the Flutter library which supports the DSL and generates the pipelines at runtime.

```

IValue preprocess(BitmapWrapper image) {
    BitmapWrapper squareImage = image.squareCrop(offset);
    BitmapWrapper postProcessed = squareImage.transform(inferenceShape.getWidth(), inferenceShape.getHeight(), 90);

    final FloatBuffer floatBuffer = Tensor
        .allocateFloatBuffer(3 * postProcessed.bitmap.getWidth() * postProcessed.bitmap.getHeight());

    final int pixelsCount = postProcessed.bitmap.getWidth() * postProcessed.bitmap.getHeight();
    postProcessed.bitmap.getPixels(buf, 0, postProcessed.bitmap.getWidth(), 0, 0, postProcessed.bitmap.getWidth(),
        postProcessed.bitmap.getHeight());

    // Torch uses CHW image format
    final int offset_g = pixelsCount;
    final int offset_b = 2 * pixelsCount;

    for (int i = 0; i < pixelsCount; i++) {
        final int c = buf[i];

        float r = ((c >> 16) & 0xff) / 255.0f;
        float g = ((c >> 8) & 0xff) / 255.0f;
        float b = ((c) & 0xff) / 255.0f;

        // imagenet normalization
        r = (r - 0.485f) / 0.229f;
        g = (g - 0.456f) / 0.224f;
        b = (b - 0.406f) / 0.225f;

        floatBuffer.put(i, r);
        floatBuffer.put(i + offset_g, g);
        floatBuffer.put(i + offset_b, b);
    }

    IValue castedInputData = IValue.from(Tensor.fromBlob(floatBuffer,
        new long[]{1, 3, inferenceShape.getWidth(), inferenceShape.getHeight()}));
    return castedInputData;
}

```

Figure 4.1: The preprocessing method in Java that Wiretronic uses for one of their models.

4.3.2.1 Domain-Specific Language

The DSL provides definitions for three different aspects of the pipeline: the model metadata, preprocessing, and postprocessing. Figure 4.2 displays the abstract syntax of the language through a metamodel, showing the main concepts of the domain and their relationships.

```

1  "model": {
2      "name": "Multihead",
3      "path": {
4          "android": "multihead.pt",
5          "ios": "multihead.mlmodel"
6      },
7      "input": {
8          "width": 380,
9          "height": 380
10     }
11 }

```

Listing 4.1: An example of how the DSL allows for specifying metadata about the model.

Using the DSL, a developer can provide metadata about the model, consisting of its name, the specific path of where to fetch the model from on iOS and Android

4. Results

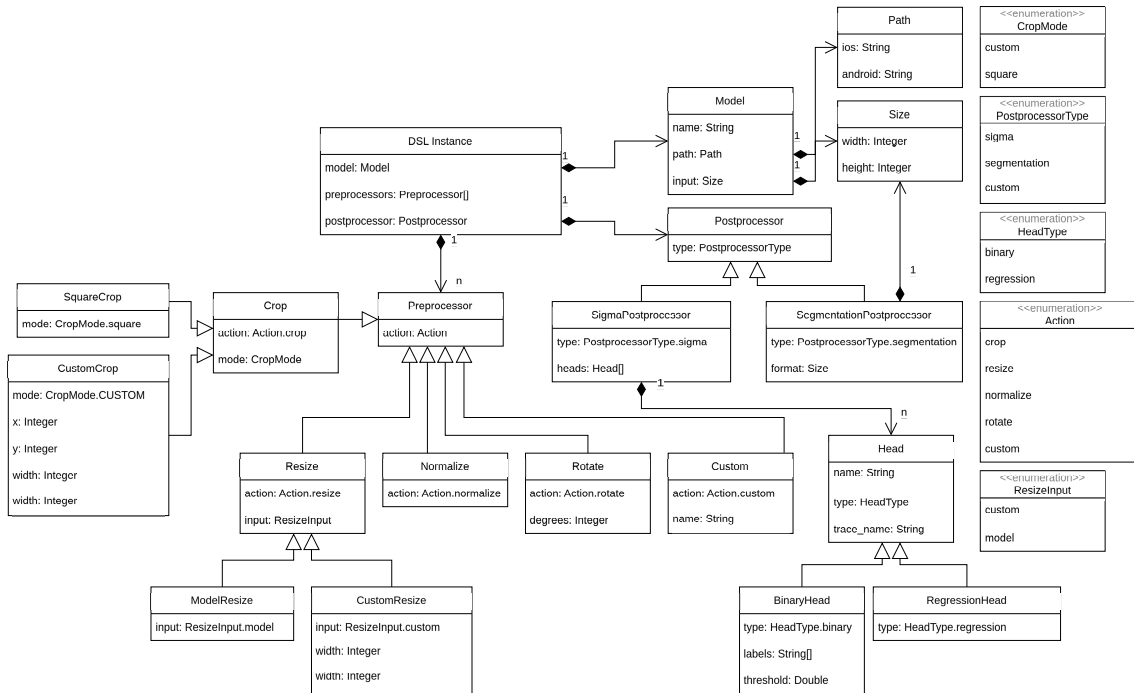


Figure 4.2: Abstract syntax of the DSL.

respectively, and the required input size of the model, which any image fed to the pipeline can be resized to. How this metadata can be defined is displayed in Listing 4.1.

Preprocessing is divided into separate steps, called preprocessors. Each preprocessor supports one specific action and can receive arguments from the developer as necessary. The DSL provides built-in support for four preprocessors: cropping, resizing, rotating, and normalizing an image. These steps are commonly used when preprocessing images for ML tasks, as the image received from e.g. the camera can be of different dimensions and orientation depending on the device configuration.

```

1  "preprocessors": [
2    {
3      "action": "crop",
4      "mode": "square"
5    },
6    {
7      "action": "resize",
8      "input": "custom",
9      "height": 380,
10     "width": 380
11   },
12   {
13     "action": "normalize"
14   }
15 ]

```

Listing 4.2: An example preprocessing configuration using the DSL.

```
1 "postprocessor": {  
2   "type": "segmentation",  
3   "format": {  
4     "height": 320,  
5     "width": 320  
6   }  
7 }
```

Listing 4.3: Example of the postprocessing in Wiretronic’s Segmentation model using our DSL.

The order of preprocessing steps is important. If, for example, an image received from the camera is 2000x2000 pixels after cropping, but the model requires an image with normalized colors of size 300x300, it would be a waste of time and computing power to apply the normalization before resizing the image, as it would require iterating through over 40 times as many pixels. Since the JSON specification does not guarantee a maintained order of object entries, the preprocessors have to be defined in an array of objects and not an object with a key for each preprocessor [51]. To accommodate this, each preprocessor is defined as an object with a key called *action* specifying the name of the step. The additional argument entries that are required for the preprocessor are then inferred by the schema through the value of the *action* key. The built-in preprocessors are defined below, and an example preprocessor configuration is displayed in Listing 4.2.

- **crop:** Allows the developer to specify a *mode*. If *mode* is *square*, it will perform a square crop in the center of the image. If *mode* is *custom*, the DSL requires the additional arguments *x*, *y*, *width*, and *height*, specified as integers.
- **resize:** Resizes the input image. The developer can choose the *input* for the measurements, if it is *custom* the image will be resized according to the arguments specified by the developer for *width* and *height*. If it is *model*, the function will use the size specified in the model metadata.
- **rotate:** Rotates an image by the number of degrees specified in the argument *degrees*.
- **normalize:** Takes no additional arguments. Normalizes the image.

While we found the preprocessing steps to be generalizable and had a large overlap in usage across models, the postprocessing was close to the opposite. Here, instead of implementing support for specific functions that can be used for many different models, we had to implement model-centric solutions.

```
1 "postprocessor": {
2   "type": "multihead",
3   "heads": [
4     {
5       "name": "is_visible",
6       "type": "binary",
7       "threshold": 0.3
8     },
9     {
10      "name": "centerpoint_x",
11      "type": "regression"
12    },
13    {
14      "name": "centerpoint_y",
15      "type": "regression"
16    }
17  ]
18 }
```

Listing 4.4: Example of the postprocessing in Wiretronic’s multi-headed model using our DSL, showcases 4 out of 11 output heads.

Listing 4.4 displays the postprocessing of Wiretronic’s multi-headed model. This figure illustrates why model output can pose a challenge when defining postprocessing of these outputs using our DSL. This model outputs 11 "heads", specific to this model. Comparing this to Listing 4.3, we see how different the two models’ outputs and postprocessing can be. During our research, we implemented functionality for these models as proof-of-concepts, displaying that the DSL can be utilized for both simple and advanced postprocessing tasks. However, if Wiretronic were to implement a completely new model, they would need to implement this in the DSL.

While extending the DSL can be suitable when introducing a new model with completely new postprocessing, there may be one-off situations where a model requires some custom operations in either the pre- or postprocessing stages. To accommodate this, we implemented a pre- and postprocessor registry, which allows developers to introduce custom functionality without the DSL being an obstacle. Contrary to Wiretronic’s current approach, where everything ML-related is handled in a library, our DSL and library would allow defining custom functionality directly in the application where it’s required. If developers then encounter the same situation in more applications, they can decide to introduce the custom step into the DSL and library permanently. The main difference between a custom implementation and an existing one is that it would require a re-release of the application since it involves writing platform-specific code that needs to be bundled with the application.

4.3.3 Model Engine

As required when opting for a model interpretation-based approach, a model engine was implemented to handle the model-to-code transformation. This model engine is displayed in Figure 4.3. When starting the application, the developer can initialize the model engine in Flutter by providing a path to the correct DSL instance.

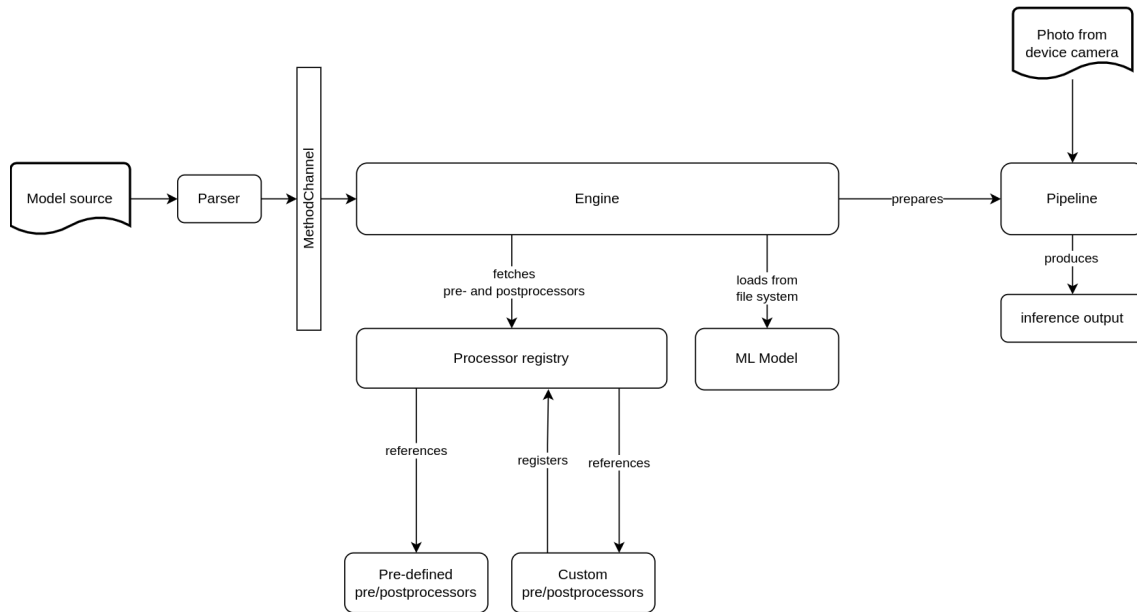


Figure 4.3: Illustration of how the model engine prepares an ML serving pipeline from a DSL instance.

This DSL instance is loaded and parsed, creating a nested dictionary referred to as the model instance. Performing the parsing in Flutter helps avoid discrepancies in parsing or file system access between platforms. After this, the model instance is fed through a MethodChannel into the platform-specific model engines. The model engine uses the model instance to fetch the correct pre- and post-processing steps for the ML model from the processor registry. Additionally, it also uses the path provided in the model instance to load the correct ML model from the file system. Upon fetching the pre- and postprocessing steps and loading the ML model, the ML serving pipeline is ready and can receive images from the device’s camera. Since the model interpretation happens at startup, any performance overheads incurred will be present on application startup and not when performing inference.

4.3.4 DSL Development Tools

We used the TypeScript tool TypeBox to develop the JSON Schema and abstract syntax that defines our DSL. TypeBox significantly reduces the amount of code having to be written compared to defining a JSON Schema manually. Additionally, it improved developer ergonomics by providing functions for set theory, allowing us to easily define complex conditional types. After the JSON Schema had been defined using TypeBox, we ran a TypeScript script that outputs the rendered JSON Schema to a JSON file. Figure 4.4 displays how TypeBox allows for separation and significantly reduced code when defining a JSON Schema, using a mock example.

```

import { Type } from "@sinclair/typebox";

const street = Type.Object({
  name: Type.String(),
  number: Type.Number({ minimum: 0 }),
});

const apartment = Type.Object({
  floor: Type.Integer({ minimum: 0 }),
  street,
});

const house = Type.Object({
  street,
});

const schema = Type.Object({
  housing: Type.Union([apartment, house]),
});

```

```

{
  "type": "object",
  "properties": {
    "housing": {
      "anyOf": [
        {
          "type": "object",
          "properties": {
            "floor": { "minimum": 0, "type": "integer" },
            "street": {
              "type": "object",
              "properties": {
                "name": { "type": "string" },
                "number": { "minimum": 0, "type": "number" }
              },
              "required": ["name", "number"]
            }
          },
          "required": ["floor", "street"]
        },
        {
          "type": "object",
          "properties": {
            "street": {
              "type": "object",
              "properties": {
                "name": { "type": "string" },
                "number": { "minimum": 0, "type": "number" }
              },
              "required": ["name", "number"]
            }
          },
          "required": ["street"]
        }
      ]
    },
    "required": ["housing"]
  }
}

```

Figure 4.4: Comparison of an example JSON schema as defined using Typebox (left) and the actual schema outputted by Typebox (right).

4.4 Artifact Evaluation

This section will go through the findings from our different evaluations of our artifact. This involves examining whether it fulfills the requirements set out at the beginning of the study along with the experiment and accompanying interviews from the third cycle.

4.4.1 Experiment Results

The results from the experiment conducted as part of our evaluation are presented here. They are presented group-wise, presenting the results from Group 1 and Group 2 for each metric. Each metric is reported per subtask.

4.4.1.1 Development Time

Overall, the artifact generated a substantial improvement in development time for all subtasks. As displayed in Table 4.1 and 4.2, this was true for both the experienced engineers (A, C) and the inexperienced engineers (B, D). Figure 4.5 displays the mean time for all participants, categorized by task and approach used. When com-

	New approach			Old approach	
	1.1	1.2	1.3	2.1	2.2
Engineer A	1	1	1	3	4
Engineer B	2	1	2	9	21

Table 4.1: The time (in minutes) it took for engineers A & B to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.

	Old approach			New approach	
	1.1	1.2	1.3	2.1	2.2
Engineer C	2	5	4	1	2
Engineer D	7	5	15	2	6

Table 4.2: The time (in minutes) it took for engineers C & D to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.

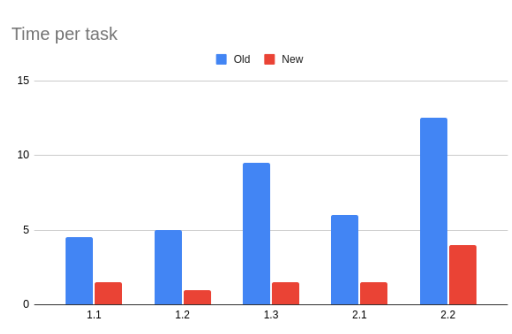


Figure 4.5: The mean time per task (in minutes) for the old and new approach, respectively.

paring an inexperienced engineer not using the artifact and one using the artifact, the average improvement in development time was 344%. If comparing experienced to inexperienced engineers before introducing the artifact, the experienced engineers on average performed 141% better than the inexperienced engineers.

		New approach			Old approach	
		1.1	1.2	1.3	2.1	2.2
Engineer A	Insertions	2	2	3	1	3
	Deletions	2	2	0	1	3
Engineer B	Insertions	2	2	4	1	6
	Deletions	2	2	0	1	4

Table 4.3: The lines of code written by engineers A & B to complete the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.

		Old approach			New approach	
		1.1	1.2	1.3	2.1	2.2
Engineer C	Insertions	2	1	4	1	1
	Deletions	2	1	0	1	0
Engineer D	Insertions	1	1	3	2	1
	Deletions	1	1	0	2	0

Table 4.4: The lines of code written by engineers C & D to complete the subtasks in the first task using the old approach, and the subtasks in the second task using the new approach.

4.4.1.2 Lines of Code

As displayed in Table 4.3 and 4.4, there were no significant differences in the absolute number of lines of code required to complete the tasks. Only in task 2.2 was there a major difference but this is attributed to normalization being a built-in function in the DSL, thus only requiring it to be enabled instead of having to perform the normalization manually. This experiment only included development on Android, however, and some of the tasks would require performing equivalent operations also on iOS, thus increasing the required lines of code when not using the DSL.

4.4.1.3 Correctness

When measuring correctness we manually tested each commit to catch any runtime failures, and statically analyzed the commits, ensuring that the commits using our artifact did not include any unnecessary code not required for the task description. As can be observed in Table 4.5 and 4.6, six subtasks implemented using the old approach were considered correct, accounting for 60%. Meanwhile, for the new approach, eight solutions were deemed correct, representing 80%.

		New approach			Old approach	
		1.1	1.2	1.3	2.1	2.2
Engineer A		correct	correct	correct	correct	incorrect
Engineer B		correct	correct	incorrect	correct	correct

Table 4.5: The correctness for engineers A & B when completing the subtasks in the first task using the new approach, and the subtasks in the second task using the old approach.

		Old approach			New approach	
		1.1	1.2	1.3	2.1	2.2
Engineer C		correct	correct	wrong	correct	correct
Engineer D		wrong	correct	wrong	wrong	correct

Table 4.6: The correctness for engineers C & D when completing the subtasks in the first task using the old approach, and the subtasks in the second task using the new approach.

4.4.2 Hypothesis Testing

We performed hypothesis testing on the development time and correctness metrics. As explained in Section 3.4.1.1, for the development time a Mann-Whitney U test was utilized, and for the correctness, a Fisher’s Exact test was used. For a test to show significance, we decided on the p-value to be lower than 0.05, as this has been the rule for strong evidence in favor of a scientific theory [52].

- Mann-Whitney U statistic: 90.0
- P-value: 0.0022820350196566465

The results of the Mann-Whitney U test indicate a statistically significant improvement in time efficiency with the new approach, as the p-value is lower than the 0.05 threshold. This finding supports discarding the null hypothesis, hence the new approach would reduce the time required to complete tasks.

- Odds Ratio: 0.375
- P-value: 0.628482972136223

The Fisher’s Exact test results in a p-value greater than 0.05, showing no statistical significance between the new and old approaches, hence keeping our null hypothesis. This result is likely due to the small sample size.

4.4.3 Requirements

Here, we evaluate the artifact with respect to the requirements defined in the first cycle. This is split into functional and non-functional requirements, and evaluated both using metrics from testing the software and subjective opinions presented by engineers during the evaluative interviews held after the experiment.

4.4.3.1 Functional Requirements

Pipeline Specification: The DSL does enable developers to specify which pre- and postprocessing steps are required for an ML model. The DSL does validate against a JSON Schema when using an IDE, both in terms of what is required for an ML model in general, and autocomplete with all pre-existing operations.

Platform-Specific Model Interpretation: The DSL does enable model interpretation in Swift and Java using the model engine illustrated in Figure 4.3, and further explained in Section 4.3.3.

Support Pre-existing and Custom Operations: Since the DSL was implemented based on the results of our interviews and repository studies, we were able to identify and implement support for the most commonly used operations, both in pre- and postprocessing. We complemented this with the previously mentioned pre- and postprocessor registry, which allows developers to include custom functionality, thus fulfilling the requirement of supporting both pre-existing and custom operations.

Support Dynamic Swapping of Configuration: The ML serving pipeline is set up through the runtime parsing of a JSON file. Thus, developers can write code that

supports changing which JSON file is loaded, and the library would instantiate a new pipeline. This is possible thanks to the model interpretation approach, performing the model-to-code transformation at runtime.

4.4.3.2 Non-functional Requirements

Usability: The goal was to make the DSL easy to learn, allowing developers to use it with a minimal training required. During the second round of interviews, the participants were asked to rate the DSL in terms of the properties of intuitiveness, learnability, and usability on a scale from 1 to 5. Intuitiveness got an average of 4.75, learnability was an average of 4.75, and usability was an average of 5. These answers indicate accomplishing our goal. Additionally, one of the inexperienced participants stated the tool provides a lower barrier of entry for contributing to the code: "I would not dare to work in this environment otherwise, using the new method makes me feel more secure" - Engineer D. However, we did get feedback on the documentation being slightly confusing, with both Engineer C & D stating that we should improve the documentation and that the large amount of text in a single place made it difficult to get an overview. Based on this feedback, we made improvements to the documentation after the interviews.

Maintainability: The main feature of the new approach is enabling easier updates of ML serving pipelines: "I think it was much better compared to without, there are so many files I don't recognize and difficult navigating the file structure" - Engineer B. The DSL enables the developers to modify the models through only one configuration file, not having to make substantial changes to the existing code.

Performance: We conducted a test measuring how long the application takes from startup to readiness, including tasks like camera setup and model loading, excluding build time. We started the application ten times using each approach, where the average of the new approach increases the startup time by 24ms. This fulfills requirement **NFR3.1**, stating that our approach should not add more than 50ms to the startup time of an application consuming the library. The test was conducted on a single computer and OS, therefore the results might differ. The full results of the trial runs are displayed in Table 4.7.

Compatibility: We manually tested the new approach across Windows, Linux, and MacOS. As long as the system had installed all the necessary software like Android Studio and Flutter, there were no issues in either system during build or runtime.

	New approach	Old approach
	1314	1381
	1323	1356
	1393	1278
	1427	1348
	1404	1340
	1392	1350
	1410	1383
	1387	1372
	1402	1374
	1371	1399
average	1382.3	1358.1
median	1392.5	1364

Table 4.7: The startup time (in milliseconds) of the application for the new and old approach respectively.

5

Discussion

This chapter discusses different ways to support several ideas provided by the interviewees during evaluation in both the first cycle and the second cycle.

5.1 Research Question 1

RQ1: How can a domain-specific language (DSL) describe an ML model (and, for example, its required inputs, outputs, and pre- and postprocessing stages)?

It is important to distinguish that our research is directed at making a DSL describing the input, output, preprocessing, and postprocessing around ML models, not the models themselves.

Through our literature review and iterative process working on the project, how an ML model can be described through a DSL depends on how generalizable you want it to be. During our research into the domain, we have found that it is quite easy to describe what happens before the data is fed into the ML model, however, the difficult part is describing what happens after. Finding a balance here was one of the more challenging tasks of the study. In the end, the choice to implement the DSL using a JSON Schema with an accompanying library allowed us to provide both a simple interface to describe ML serving pipelines and a way to implement custom, one-off features without slowing down development.

Our approach allows developers to specify metadata about the model, consisting of its name, path on the device, and input shape. The preprocessing is described as a series of steps, where we implemented support for the most common actions used in Wiretronic's current ML serving pipelines in addition to the possibility of defining custom steps. Lastly, the DSL supports specifying the required postprocessing actions. While the preprocessing actions were found to be quite trivial and generalizable, the postprocessing steps are usually different between each model, opting for having to implement custom functionality to handle the model outputs. Here, the functionality to easily be able to define custom postprocessing actions is necessary.

Summary

Representing an ML model's input, output, preprocessing, and postprocessing steps in a JSON configuration file requires a minimum of model metadata, preprocessing actions, and postprocessing actions. By utilizing JSON Schemas you can easily expand upon what features to have depending on the situation.

5.2 Research Question 2

RQ2: How can we best implement and utilize the DSL in a concrete setting, specifically in the development of cross-platform mobile applications?

One aim with the DSL was to create a unified interface for not only multiple platforms but also for engineers of different backgrounds. This meant that we did not want to make it overly related to the underlying platforms, since this could cause confusion or unfamiliarity for ML-focused engineers. Furthermore, we did not want to make the DSL too restricting, offering experienced engineers the possibility to combine the DSL with custom, platform-specific functionality.

While the DSL is aimed at cross-platform mobile development, we did not want to make it tied to the technique currently used at Wiretronic, for example as an internal DSL written in Dart (for Flutter). This connects to the previously mentioned point of creating a unified interface across platforms and experience levels, but it also allows for porting or extending the DSL to additional platforms. We developed the DSL and accompanying library so that if Wiretronic decides to shift its cross-platform development to another technique, the DSL would not require any modifications.

To accommodate the initial requests made by engineers to not require a complete re-release of the application upon changes to the ML serving pipeline, we implemented the DSL using a model interpretation approach, instead of using code generation. Any application consuming our accompanying library could fetch a remote file written in our DSL and dynamically set up the ML serving pipeline, without having to re-release the application.

Summary

The DSL and accompanying library were implemented to support different underlying techniques, engineers of different backgrounds, and making changes to the ML serving pipeline without re-releasing the application.

5.3 Research Question 3

RQ3: To what extent does the introduction of a DSL and an accompanying library improve the developer experience in the aspects of maintenance, feature development, time-saving, and resource planning?

From our controlled experiment and two rounds of interviews, it has been made clear that a DSL designed in a familiar format can aid in lowering the barrier of entry in this area of development. This may be the largest improvement when comparing the previously used approach, as new engineers can contribute and experiment in development. Through both objective and subjective metrics, our evaluation showed that the engineers worked faster and more confidently while using our approach, partly thanks to the ML-related functionality being isolated into a single file and format. In addition, the results in Section 4.4.1 show us improvement in all metrics using our new approach. The average improvement in development time was 344%, which is also backed by the hypothesis testing performed in Section 4.4.2. The correctness improved by 20% in the controlled experiment, but we could not prove statistical significance. The engineers did however state in the interviews following the experiment they still felt more secure using the new approach. If the DSL can help more engineers contribute to this area of development it can help in all the aspects stated in this research question. More engineers will be able to perform maintenance tasks and develop new features, further helping Wiretronic deliver features faster and easing their planning.

Summary

The DSL does lower the complexity of edge-deployed ML at Wiretronic. The DSL and accompanying library make the entry into the field quicker, while also enabling the engineers to do the work faster and more confidently. Hence, improving the developer experience in the aspects outlined in the research question.

5.4 Cross-Platform Communication

Due to the nature of cross-platform development, there will be many instances of communication between the Flutter layer and the native layer through MethodChannels. During the development of our library, we ran into many instances of having to debug on both sides of the MethodChannels. This can become a very tedious and time-consuming task, and with our library solving the issues of layer communication, we can effectively ease the need for debugging for the developers.

In the future, we may see a shift away from using channels and data serialization for inter-layer communication. React Native explores this in their new architecture, which is under development at the time of writing. Here, the native code is written in C++ and the cross-platform layer (in this case, JavaScript) can hold references to C++ objects and vice-versa, calling functions directly on these objects [12].

5.5 Threats to Validity

In our project and research, we have three main threats to validity.

5.5.1 Internal Validity

Internal validity is of concern when we examine causal relations [53]. We aimed to ensure internal validity in our study by using a controlled setting, with the intent of eliminating any confounding factors. The elicitation of our requirements and scope was defined through only two interviews and casual conversations, each conducted with the one mainly responsible for the Android part, and the other responsible for the iOS part.

We evaluated our DSL and library with the developers at Wiretronic as the problem and research area were brought to light by them. This is a small group, and this may cause problems with internal validity. If creating random groups for the experiment, there is a chance that uneven levels of previous experiences in the area become a confounding factor. As a mitigation to this issue, we used stratified sampling and a cross-over design, in which participants from both groups are exposed to both treatments, thereby turning each participant into their own control. As part of the evaluation we also used manual static testing for the correctness, which may introduce human error or bias. To mitigate this we also ran the code to see if the solutions would introduce any runtime errors.

The two experienced participants in the experiment were also the interviewees in the initial interviews, which might introduce bias or influence their performance. This overlap could affect the way they approached the tasks. We are aware of this as a potential threat but tried to minimize its impact through clear communication.

Another factor that may affect internal validity is the possibility that participants may not be honest in their feedback. One of us has previously been employed at Wiretronic, so the developers might provide positive feedback to help us in our work. To mitigate this issue, we will clearly communicate that the respondents should give answers and prefer tasks as they naturally would. Additionally, having been colleagues in the past, we know the employees well enough that they will be comfortable giving honest critiques.

During the experiment, one participant performed worse than the others, which can be considered an outlier. We ran the statistical analysis again without the data from this engineer, but we still observed statistical significance. Mitigating the impact on our results.

5.5.2 External Validity

External validity is to what extent it is possible to generalize the findings [53]. Our DSL will be designed based on current literature and studies, however our proof-of-concept library will be designed and implemented in only the single ecosystem of Wiretronic. This population comprises of proficient software engineers with several

years of experience. Thus, the homogeneity and small size of our experiment may not be generalizable.

Furthermore, as our scope limits us to deploying ML models on edge devices, this study may not be applicable to all sorts of models or devices. Additionally, integration and compatibility with existing ML frameworks or external platforms may pose a challenge as the DSL will be tailored to the needs of Wiretronic.

5.5.3 Construct Validity

Construct validity reflects if the measurements really represent what they are meant to do [53]. During the controlled experiment we saw the low amount of lines of code needed for our new DSL approach. It can be argued that our experiment is not comprehensive enough, but the tasks were designed with this specifically in mind. One of the main pain points of Wiretronic was the need to write unnecessary amounts of code to make these small tweaks. Another point is the completeness of the DSL, the experiment does not capture all functionalities of the DSL.

6

Conclusion

In this thesis, we have developed a domain-specific language specifically for defining an ML serving pipeline in a cross-platform environment. The study aimed to simplify the deployment and maintenance of ML models on edge devices, focusing on the unique challenges present in this environment. These challenges included software deployment to external devices and working in cross-platform environments with several code bases.

The first cycle was about defining the scope through initial interviews with Wiretronic engineers and obtaining an overview of the domain through researching relevant literature. Combining Wiretronic’s needs and thoughts with what technologies would be the optimal approach for them, ultimately comes down to the lightweight and easy-to-learn JSON Schema approach. Allowing the developers to in one single configuration file change quickly what before would need a redeployment of their library. The second cycle used the information gathered in the first cycle to design and implement the artifact. Using the guidance from the research, interviews, and engineers to ensure usefulness for the company and compatibility with their current systems. Lastly, the third cycle conducts the controlled experiment followed by second interviews to evaluate the artifact, both in terms of metrics and subjective reactions from the engineers.

Our research confirms that a DSL can reduce the complexity typically involved when writing equivalent functionality for several platforms. By abstracting platform-specific code, the DSL enables developers to define ML models and their required inputs, outputs, and pre- and postprocessing stages, simplifying the aspects mentioned in **RQ3**.

In conclusion, the development of our artifact represents a step toward simplifying the deployment of ML models especially on smartphones in a cross-platform environment. It bridges the gap between ML engineers and other developers, enabling common understanding and lowering the barrier of entry to cross-platform development at Wiretronic.

6.1 Limitations

Our artifact currently only has the most trivial iOS files implemented. Due to the lack of access to a Mac and an Apple developer account, it was not possible to build, test, and run the code, making that implementation difficult. With access

to a Mac, it should be quite easy to finalize the iOS version of the artifact, as the cross-platform functionality already exists and will stay the same.

6.2 Future Work

The work done during this study can be expanded on, both within the domains covered and across other areas of software engineering.

6.2.1 DSL Expansion

There are numerous possibilities for adding more features to the DSL, increasing its functionality and flexibility. For instance, a scenario where Wiretronic needs to execute different pre- and postprocessing steps depending on the deployment platform. Additionally, the current requirement of having a separate configuration file for each model could be streamlined by extending the DSL to integrate all model configurations into a single, more compact file. While there are many possibilities, it is important to note this could also lead to increased complexity.

6.2.2 Deeper Exploration Within the Domain

If exploring further applicability of the DSL within the same domain, it can be beneficial to implement a simple way to specify if pipeline steps should be executed on only some platforms. Furthermore, to continue the work within the same or adjacent domains, a more generalizable solution can be explored. This study focused only on the needs of a single company, primarily developing ML models for image-based tasks. Due to the small size of Wiretronic's AI division, it would be valuable to study the applicability of similar artifacts in a larger organization. This could both help in refining and generalizing the DSL and help combat the validity threats related to the small sample size of this study.

Expanding the focus to adjacent domains, it can be beneficial to study both the potential of using a lightweight DSL to describe ML serving pipelines in environments other than smartphones and to explore its applicability for scenarios other than image-based ML.

6.2.3 Expansion to Other Domains

Although the developed artifact was tailored to the needs of Wiretronic, future research can be conducted into the impact of employing lightweight DSLs in other settings, for organization-specific scenarios. While the development of a more traditional DSL can be expensive in both development and adoption time, opting for a lightweight and familiar approach can lower both development and adoption time, while simultaneously lowering the barrier of entry for contributing to the DSL.

Bibliography

- [1] F. Chollet, “Deep learning with python,” in 2nd ed. Manning Publications, Oct. 2021, ch. 1.2.7.
- [2] *wirevision*. [Online]. Available: <https://apps.apple.com/se/app/wirevision/id1543684933>.
- [3] *Husqvarna Gear Identifier*. [Online]. Available: <https://apps.apple.com/se/app/husqvarna-gear-identifier/id6464394076>.
- [4] *Flutter*. [Online]. Available: <https://flutter.dev/>.
- [5] *Core ml*, Accessed: 2024-01-22, 2023. [Online]. Available: <https://web.archive.org/web/20231126225328/https://developer.apple.com/machine-learning/core-ml/>.
- [6] *Custom models / ml kit*, Accessed: 2024-01-22, 2023. [Online]. Available: <https://web.archive.org/web/20231208182614/https://developers.google.com/ml-kit/custom-models>.
- [7] L. Lai and N. Suda, *Rethinking machine learning development and deployment for edge devices*, 2018. arXiv: 1806.07846 [cs.LG].
- [8] S. P. Bayerl, T. Frassetto, P. Jauernig, *et al.*, “Offline Model Guard: Secure and Private ML on Mobile Devices,” in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2020, pp. 460–465. DOI: 10.23919/DATE48585.2020.9116560.
- [9] H. Song, R. Dautov, N. Ferry, A. Solberg, and F. Fleurey, “Model-based fleet deployment of edge computing applications,” Association for Computing Machinery, Inc, Oct. 2020, pp. 132–142, ISBN: 9781450370196. DOI: 10.1145/3365438.3410951.
- [10] Q. Zhang, X. Li, X. Che, *et al.*, “A Comprehensive Benchmark of Deep Learning Libraries on Mobile Devices,” in *Proceedings of the ACM Web Conference 2022 (WWW ’22)*, Virtual Event, Lyon, France: ACM, Apr. 2022, p. 10. DOI: 10.1145/3485447.3512148.
- [11] *Stack overflow developer survey 2023*, Accessed: 2024-01-22, 2023. [Online]. Available: <https://web.archive.org/web/20240121213425/https://survey.stackoverflow.co/2023/#most-popular-technologies-misc-tech-prof>.

- [12] React Native, *Why a New Architecture*, Accessed: 2024-02-22, 2023. [Online]. Available: <https://web.archive.org/web/20231206110642/https://reactnative.dev/docs/the-new-architecture/why>.
- [13] Flutter, *Writing custom platform-specific code*, Accessed: 2024-02-22. [Online]. Available: <https://web.archive.org/web/20240214020240/https://docs.flutter.dev/platform-integration/platform-channels>.
- [14] M. Brambilla, J. Cabot, and M. Wimmer, *Model-Driven Software Engineering in Practice*. Springer International Publishing, 2017, ISBN: 978-3-031-01421-5. DOI: 10.1007/978-3-031-02549-5. [Online]. Available: <https://link.springer.com/10.1007/978-3-031-02549-5>.
- [15] E. de Araújo Silva, E. Valentin, J. R. H. Carvalho, and R. da Silva Barreto, *A survey of model driven engineering in robotics*, Feb. 2021. DOI: 10.1016/j.cola.2020.101021.
- [16] B. Lelandais, M. P. Oudot, and B. Combemale, “Applying model-driven engineering to high-performance computing: Experience report, lessons learned, and remaining challenges,” *Journal of Computer Languages*, vol. 55, Dec. 2019, ISSN: 25901184. DOI: 10.1016/j.cola.2019.100919.
- [17] A. Moin, M. Challenger, A. Badii, and S. Gunnemann, “Supporting ai engineering on the iot edge through model-driven tinyml,” Institute of Electrical and Electronics Engineers Inc., 2022, pp. 884–893, ISBN: 9781665488105. DOI: 10.1109/COMPSAC54236.2022.00140.
- [18] S. Vaupel, G. Taentzer, J. P. Harries, R. Stroh, R. Gerlach, and M. Guckert, *Model-driven development of mobile applications allowing role-driven variants*, 2014.
- [19] S. Vaupel, D. Strüber, F. Rieger, and G. Taentzer, “Agile bottom-up development of domain-specific ides for model-driven development.,” in *FlexMDE@MoDELS*, 2015, pp. 12–21.
- [20] S. P. Jácome-Guerrero, M. Ferreira, and A. Corral, “Software development tools in model-driven engineering,” in *2017 5th International Conference in Software Engineering Research and Innovation (CONISOFT)*, 2017, pp. 140–148. DOI: 10.1109/CONISOFT.2017.00024.
- [21] U. Frank, “Domain-specific modeling languages: Requirements analysis and design guidelines,” in May 2013, pp. 133–157, ISBN: 978-3-642-36653-6. DOI: doi:10.1007/978-3-642-36654-3_6.
- [22] R. F. Paige, D. S. Kolovos, and F. A. Polack, “A tutorial on metamodelling for grammar researchers,” vol. 96, Elsevier B.V., Dec. 2014, pp. 396–416. DOI: doi:10.1016/j.scico.2014.05.007.
- [23] M. Mernik, J. H. Cwi, and A. M. Sloane, *When and how to develop domain-specific languages*, 2005.
- [24] *Xtext*, Accessed: 24.01.24. [Online]. Available: <https://eclipse.dev/Xtext/>.

-
- [25] A. Iung, J. Carbonell, L. Marchezan, *et al.*, “Systematic mapping study on domain-specific language development tools,” *Empirical Software Engineering*, vol. 25, pp. 4205–4249, 5 Sep. 2020, ISSN: 15737616. DOI: 10.1007/s10664-020-09872-1.
- [26] Z. M. Korani, A. Moin, A. R. da Silva, and J. C. Ferreira, *Model-driven engineering techniques and tools for machine learning-enabled iot applications: A scoping review*, Feb. 2023. DOI: 10.3390/s23031458.
- [27] A. Moin, M. Challenger, A. Badii, and S. Günemann, “A model-driven approach to machine learning and software modeling for the IoT,” *Software and Systems Modeling (SoSyM)*, 2022. DOI: 10.1007/s10270-021-00967-x. [Online]. Available: <https://doi.org/10.1007/s10270-021-00967-x>.
- [28] A. Moin, *ML-Quadrat*, <https://github.com/arminmoin/ML-Quadrat>, Accessed: 2024-01-24, 2024.
- [29] *Domain-specific languages, projectional editor and mps*, Accessed: 2024-02-26, 2023. [Online]. Available: <https://www.jetbrains.com/help/mps/mps-faq.html#dsl-projectional-editor-mps>.
- [30] J. Zhao, T. Tiplea, R. Mortier, J. Crowcroft, and L. Wang, “Data analytics service composition and deployment on edge devices,” Association for Computing Machinery, Inc, Aug. 2018, pp. 27–32, ISBN: 9781450359047. DOI: 10.1145/3229607.3229615.
- [31] S. Idowu, D. Strüber, and T. Berger, “Emmm: A unified meta-model for tracking machine learning experiments,” in *2022 48th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, 2022, pp. 48–55. DOI: 10.1109/SEAA56994.2022.00016.
- [32] J. Giner-Miguel, A. Gómez, and J. Cabot, “Describeml: A tool for describing machine learning datasets,” in *Proceedings of the 25th International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings*, 2022, pp. 22–26.
- [33] E. Chavarriaga, F. Jurado, and F. D. Rodríguez, “An approach to build json-based domain specific languages solutions for web applications,” *Journal of Computer Languages*, vol. 75, p. 101203, 2023, ISSN: 2590-1184. DOI: <https://doi.org/10.1016/j.cola.2023.101203>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2590118423000138>.
- [34] A. Habib, A. Shinnar, M. Hirzel, and M. Pradel, “Finding data compatibility bugs with json subschema checking,” Association for Computing Machinery, Inc, Jul. 2021, pp. 620–632, ISBN: 9781450384599. DOI: 10.1145/3460319.3464796.
- [35] LinkedIn, *What are the most common data conversion formats and standards in your industry?* Accessed: 2024-02-19, 2023. [Online]. Available: <https://www.linkedin.com/advice/0/what-most-common-data-conversion-formats-standards>.

- [36] “The top 10 data interchange or data exchange format used today.” Accessed: 2024-02-19. (May 2023), [Online]. Available: <https://aster.cloud/2023/05/11/the-top-10-data-interchange-or-data-exchange-format-used-today/>.
- [37] JSON Schema, *Getting started step by step*, Accessed: 2024-02-19, 2023. [Online]. Available: <https://json-schema.org/learn/getting-started-step-by-step>.
- [38] L. Attouche, M.-A. Baazizi, D. Colazzo, G. Ghelli, C. Sartiani, and S. Scherzinger, *Validation of modern json schema: Formalization and complexity*, 2024. arXiv: 2307.10034 [cs.DB].
- [39] E. Knauss, “Constructive master’s thesis work in industry: Guidelines for applying design science research,” Dec. 2020. [Online]. Available: <http://arxiv.org/abs/2012.04966>.
- [40] K. Matzler and H. Hinterhuber, “The kano model: How to delight your customers,” Feb. 1996, pp. 313–327. [Online]. Available: <https://www.researchgate.net/publication/240462191>.
- [41] R. C. Martin, “Design principles and design patterns,” in 2000, pp. 12–14.
- [42] L. Gao, “Latin squares in experimental design,” *Michigan State University*, 2005.
- [43] S. Baltés and P. Ralph, “Sampling in software engineering research: A critical review and guidelines,” *Empirical Software Engineering*, vol. 27, no. 4, p. 94, 2022.
- [44] K. J. Stol and B. Fitzgerald, “The abc of software engineering research,” *ACM Transactions on Software Engineering and Methodology*, vol. 27, 3 Sep. 2018, ISSN: 15577392. DOI: 10.1145/3241743.
- [45] SciPy Developers, *Scipy.stats.mannwhitneyu*, Accessed: 2024-05-20. [Online]. Available: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.mannwhitneyu.html>.
- [46] SciPy Developers, *Scipy.stats.fisher_exact*, Accessed: 2024-05-20. [Online]. Available: https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.fisher_exact.html.
- [47] Interaction Design Foundation - IxDF, *What are user stories?* <https://www.interaction-design.org/literature/topics/user-stories>, Accessed: 02.02.24, May 2016.
- [48] S. Lauesen, *Software Requirements-Styles and Techniques*. Jan. 2002.
- [49] J. Huang, Y. F. Li, and M. Xie, “An empirical analysis of data preprocessing for machine learning-based software cost estimation,” vol. 67, Elsevier B.V., Nov. 2015, pp. 108–127. DOI: 10.1016/j.infsof.2015.07.004.
- [50] *Non-functional requirements: Examples, types, and how to approach them*, <https://www.altexsoft.com/blog/non-functional-requirements/>, Accessed: 2024-02-26, 2023.

- [51] *The json data interchange syntax*, Accessed: 2024-04-14, 2017. [Online]. Available: <https://web.archive.org/web/20240407090452/https://ecma-international.org/publications-and-standards/standards/ecma-404/>.
- [52] B. B. McShane, D. Gal, A. Gelman, C. Robert, and J. L. Tackett, “Abandon statistical significance,” *The American Statistician*, vol. 73, no. sup1, pp. 235–245, 2019. [Online]. Available: <https://doi.org/10.1080/00031305.2018.1527253>.
- [53] C. Wohlin, P. Runeson, M. Host, M. C. Ohlsson, B. Regnell, and A. Wesslen, *Experimentation in software engineering*, 2012. [Online]. Available: <https://link.springer.com/book/10.1007/978-3-642-29044-2>.

A

Appendix 1 - Initial Interviews

- Can you explain what code is equivalent but different in implementation across platforms, and give examples from your experience where you had to write or manage it?
- How often do you have to edit this code?
- How much work is it to make adjustments to this code?
- Is the file structure consistent across other projects you work on?
 - Do you have a standardized file structure or guideline to use across projects?
- Is the code in its current form reusable across other repositories, or is it context-specific?
- What IDE do you use for this work?
 - Do you use different IDEs for the different languages (Java/Swift)?
- How do you do testing on native mobile code?
- Would it speed up your workflow if you had a tool to generate platform-specific code for you by having to specify it only once in a file, similar to a configuration file?
 - If so, how would you imagine the tool being implemented?

B

Appendix 2 - Experiment Interviews

- How extensive is your previous experience in deploying ML models in a cross-platform mobile environment on a scale from 1 (no experience at all) to 5 (very experienced)?
 - How many years have you spent doing that work?
- What was your general impression of using the DSL to implement and make changes to an ML serving pipeline?
- After using the DSL in the experiment, how would you rate it for the following properties (from 1 to 5):
 - Intuitiveness - how intuitive and natural was the DSL to use?
 - Learnability - how fast could you get up to speed using the DSL?
 - Usability - how usable was the DSL for the task at hand?
- Based on your own experience in implementing ML models: How would you rate the usefulness of the DSL and the accompanying library on a scale from 1 (not useful at all) to 5 (very useful) in your day-to-day development workflow, and why?
- Can you think of any scenarios where the DSL in its current form would be particularly useful?
 - If so, are there any features you would want the DSL to have?
- Can you think of any scenarios where the DSL in its current form would not be useful?
 - If so, are there any existing features that are unnecessary?

C

Appendix 3 - Mann-Whitney U Test Code

The following Python code was used to perform the Mann-Whitney U test on the development time data for the old and new approaches:

```
import scipy.stats as stats

# Old approach times
old_approach_times = [2, 3, 4, 4, 9, 21, 2, 5, 7, 15]
old_approach_times.sort()

# New approach times
new_approach_times = [1, 1, 1, 2, 2, 1, 2, 1, 2, 6]
new_approach_times.sort()

# Perform the Mann-Whitney U test
u_statistic, p_value =
    stats.mannwhitneyu(old_approach_times, new_approach_times)

print(f"Mann-Whitney U statistic: {u_statistic}")
print(f"P-value: {p_value}")
```


D

Appendix 4 - Fisher's Exact Test Code

The following Python code was used to perform the Fisher's Exact Test on the correctness data for the old and new approaches:

```
import scipy.stats as stats

# Old approach: [correct, incorrect]
# New approach: [correct, incorrect]
old_approach = [6, 4]
new_approach = [8, 2]

# Create a contingency table
contingency_table = [old_approach, new_approach]

# Perform Fisher's Exact Test
oddsratio, p_value = stats.fisher_exact(contingency_table)

print(f'Odds Ratio: {oddsratio}')
print(f'P-value: {p_value}')
```


E

Appendix 5 - Artifact Code

The code and artifact for this thesis are available in the following Google Disk folder. Note that it is not fully complete as Wiretronic did not wish for everything to be published.

[Click here to go to the Google Disk folder](#)