Universität Zürich[UZH]

Geographisches Institut

MyPlaces – An Adaptive Mobile Map Application

GEO 511 Master's Thesis

Author: Johannes Guler, 19-756-741

Supervised by: Dr. Tumasch Reichenbacher

Faculty representative: Prof. Dr. Sara Irina Fabrikant

30.09.2025

Abstract

Mobile map applications have become crucial tools for spatial navigation and decision-making, yet they typically present static interfaces and follow a one-size-fits-all approach. Consequently, the user is overwhelmed with irrelevant information. This thesis presents MyPlaces, an adaptive mobile map application that dynamically adjusts the visualisation of point of interests (POI) and basemap based on real-time geographic relevance (GR) assessments. By implementing a slightly modified version of De Sabbata's (2013) GR framework, through a novel dual-model machine learning architecture, the system processes contextual variables including spatio-temporal, semantic, behavioural, and geographic environment factors to generate personalised relevance scores.

The prototype is developed for iOS using Swift as the main programming language and the ArcGIS Maps SDK. Additionally, XGBoost models are employed for thematic activity prediction and relevance scoring, achieving a dynamic performance and quick processing time. The visual adaptation strategy makes use of cartographic design principles through size, opacity, and colour saturation encoding, while an intelligent aggregation algorithm reduces visual complexity depending on relevance and map scale. Technical evaluation across urban, rural, and outdoor environments demonstrates the system's ability to process thousands of POIs with only short loading times. Privacy is granted through the usage of the device's core data and local relevance processing.

The modular, open human-computer interaction source architecture provides a research platform for further cartographic research in the adaptivity and mobile map design domain. While limitations include reliance on synthetic training data and absence of user evaluation studies, MyPlaces successfully proves that GR-based adaptive mapping is technically feasible in real-time mobile environments, enabling baselines and suitable design patterns for future research in intelligent spatial information systems.

Keywords**:** adaptive mobile mapping, context-aware computing, geographic relevance, human-computer interaction, location-based services, machine learning, mobile map design, Swift, iOS, ArcGIS Maps SDK, ArcGIS Location Platform

Acknowledgements

I wish to express my gratitude to the following:

Dr Tumasch Reichenbacher for his valuable guidance through this master's thesis and taking his time to give constructive feedback. Insightful discussions with him were instrumental to the development of this project

Dr Peter Ranacher who provided helpful feedback on the initial concept, which provided a strong foundation for the work that followed

Prof. Dr Ross Purves for feedback on the model training process and the architectural design of the application

Matthias Schenker and Philipp Marty from Esri Schweiz for their technical support and for providing access to the Esri Location Platform and the ArcGIS licenses

My parents and friends for proofreading the thesis and offering valuable feedback on both form and content.

Table of Contents

List of Figures

Tables

Abbreviations

AMMA	Adaptive Mobile Map Application
API	Application programming interface
BFS	Bundesamt für Statistik (Switzerland)
CPA	Conjunctive partial absorption
CPL	Continuous preference logic
CV	Content view
CVM	Content view model
GCD	Generalised conjunction or disjunction
GIN	Geographic information need
GIO	Geographic information object
GIVA	Geographic Information Visualisation and Analysis group, University of Zurich
GPS	Global positioning system
GR	Geographic relevance
HCI	Human-computer interaction
IR	Information retrieval
MAE	Mean absolute error
ML	Machine learning
MSE	Mean square error
OOP	Object-oriented programming
OSM	OpenStreetMap
POI	Point of interest
POIM	POI model
RM	Relevance model
RMSE	Root mean square error
RQ	Research question
SDK	Software Development Kit
TM	Thematic model
UI	User interface

1 Introduction

Picture yourself arriving in Zürich Hauptbahnhof during evening rush hour, smartphone in your hand and searching for a place to eat. Your map application displays various kinds of places, including shops, museums, or restaurants that have already closed. The cognitive burden of filtering through this visual noise while navigating through the streets demonstrates a fundamental challenge in modern cartography: How can mobile maps intelligently adapt to show the information needed when needed?

1.1 The Evolution of Digital Mapping

Cartography began with simple paper maps and has undergone two major paradigm shifts: from static maps to interactive web-based maps, and then to mobile maps that are embedded in everyday life (Clarke, 2001; Harrell, 2025). Those transitions have changed the way maps are accessed, used, and perceived. Maps have shifted from fixed representations to interactive interfaces, and ultimately to ubiquitous spatial tools that are embedded in every smartphone (Reichenbacher, 2004).

The first shift occurred with the digitalisation of traditional cartography practices. Following the emergence of web mapping in the early 2000s,static paper maps that were printed in atlases or charts were published online. This advance provided the tools for map interactions, such as panning or zooming, but the map content itself remained the same for all users (Been, 2002).

A second transition began with the technological advances of smartphones. These mobile devices, equipped with global positioning system (GPS) and internet connections, can turn maps into responsive, context-aware tools. Access to maps was no longer bound to fixed locations and included real-time positioning capabilities, providing users with essential navigational and spatial awareness tools. As a result, maps became increasingly integrated into daily life.

Despite technological advances, mobile map applications continue to face limitations, with the most significant being screen size and limited adaptation to users' needs (Harrie, Sarjakoski, & Lehto, 2002). As Reichenbacher has observed, 'The geovisualisation on small displays is dominated by the constraint of the small display' (2004, p.4), which in turn forces the map content to be reduced and generalised. Various platforms have introduced dynamic and real-time features. However, they remain committed to the idea that one map should serve all users equally. Although the technology has advanced, the conceptual model of how maps communicate information has remained largely static (Bartling, Reichenbacher, & Fabrikant, 2023; Reichenbacher, 2004). The resulting gap serves as the foundation for the next evolution needed for mapping, which is adaptivity. As several scientific studies have argued, genuinely user-centred maps should go beyond interactivity and react to contextual changes (Bartling et al., 2023; Reichenbacher, 2004, 2019). The adaptation should change what is shown and how it is shown, based on the user of the map, time, and location. Efficiently communicating the needed information to the user would help maximise the usefulness of the constrained screen size.

1.2 Research Objectives

The goal of this thesis is to explore how real-time adaptive mapping on mobile devices can be developed and implemented to improve the users' ability to discover and interact with geographically relevant information. As mobile users frequently rely on maps in unfamiliar environments, maps must not only display accurate information but also adapt to the user's preferences and context. Traditional static maps often fail to meet these dynamic needs. Therefore, this research investigates how adaptive mobile maps can be designed to respond in real time to changing user contexts and deliver a more personalised and valuable experience.

The primary objective of this thesis is to prove the technical feasibility of real-time adaptive mapping on mobile devices. A functional prototype is developed to demonstrate the integration of adaptive features, such as dynamic point of interest (POI) filtering or context-sensitive visual adjustments. The secondary objective is to establish design patterns and best practices for adaptive map interfaces by drawing from current literature in geovisualisation and human-computer interaction (HCI). The implementation tests hypotheses around user interaction, accessibility, and information prioritisation based on geographic relevance (GR). The tertiary objective is to contribute openly to the mapping and developer community by making the research findings, methodology, and prototype codebase publicly available. The thesis aims to inspire further innovation and adoption of adaptive mapping in applications.

To support these objectives, the following research questions are addressed:

- Research Question 1: How can geographic relevance be modelled with user context and environmental factors to enable personalised POI ranking in a mobile map application?
- Research Question 2: How can a mobile map application be designed and implemented to dynamically adapt map content and visualisation based on real-time geographic relevance?
- Research Question 3: To what extent does the developed prototype demonstrate the feasibility and potential of real-time adaptive mapping as a foundation for future research in adaptive maps and cartographic design?

Answering RQ1 involves synthesising existing scientific knowledge to identify relevant factors of user context and study already proposed adaptive mechanisms. Based on those mechanisms, the modelling of GR needs to be conducted with a scoring logic.

To answer RQ2, the focus is set on HCI, design principles, and technical implementation to make adaptation time efficient and computationally efficient. A working prototype called 'MyPlaces' is developed in Swift for the iOS ecosystem. The conceptualisation and implementation of MyPlaces is guided by the following central use cases: on-the-go search and location exploration. These situations should reflect the core aims of adaptive maps: to help users quickly and confidently locate places of high relevance to their current needs while minimising cognitive effort.

For RQ3, the prototype is evaluated for technical feasibility by assessing the speed and accuracy of adaptations and predictions using synthetic test data. The implementation is structured in a modular and open-source manner to facilitate extensibility, enabling further development and experimentation by researchers. In particular, the system is intended to serve as a future research tool for the Geographic Information Visualisation and Analysis (GIVA) group at the University of Zurich, supporting ongoing investigations into adaptive cartography and user-centred geographic relevance modelling.

1.3 Outline

The following chapter lays the theoretical foundations for this research topic, beginning with significant developments conducted in the field of digital mapping. The concept of GR is assessed in detail, followed by an examination of mobile design principles for mobile applications. Building on this, technical implementation possibilities for an adaptive mobile map application are discussed. Chapter 3 covers the methodology of the implementation of MyPlaces, including the data used, the modelling approaches, as well as the technical implementation and design process of the adaptive mobile map application (AMMA) prototype. Functional, technical and scientific use cases are defined. The resulting application with its adaptive features as well as the results of the qualitative and quantitative assessment metrics are presented in Chapter 4. Chapter 5 discusses the results in relation to this thesis's research questions, use cases, and the theoretical background. Key insights and the application's limitations are highlighted, and opportunities for further improvement and research are identified. Finally, Chapter 6 concludes with a summary of the main findings and contributions.

2 Theoretical Background

The foundation of this thesis is based on interdisciplinary research that informs the design and evaluation of an AMMA. The following sections explore the most relevant scientific findings and introduce conceptual frameworks that inform the development process of a mobile map application. The chapter is divided into three main parts. First, the concept of GR, along with its assessment criteria, are introduced. Second, the essential design principles from cartography, HCI, and information science are presented. Those two areas of research form the base for the implementation of the AMMA prototype. Lastly, technical findings of ideal implementations are discussed.

2.1 Geographical Relevance

To understand user needs and determine the relevance of information for a user, this section introduces GR as the primary theoretical foundation that helps create an adaptive map design. It first defines the user context and the factors that drive it, forming the groundwork for modelling personalised relevance. Building on this. the conceptual framework of GR is presented. The section concludes with the examination of key factors influencing GR, as well as their varying importance for different situations.

2.1.1 The Concept of Geographic Relevance

Figuring out the relevance of information for a user is one of the most crucial parts of adaptive mapping. Although it may seem intuitive, the assessment can be rather complex. Depending on the research field, definitions and approaches can vary (Reichenbacher, 2007).

Early text-based research in information retrieval (IR) science already acknowledged the importance of context for assessing relevance (Wilson, 1973). Identifying relevant information requires considering not only objective relevance but also the users subjective perspective. A place’s relevance therefore emerges only in alignment with the user’s context.

GR extends the classical concept of IR to spatial decision-making, focusing not on metadata or documents but rather on the entities in the real world (Reichenbacher & De Sabbata, 2011; Zingaro & Reichenbacher, 2022b). Therefore, the significance of entities, such as places, events, or features, is assessed within a given user context combination that incorporates both spatial and temporal components (De Sabbata & Reichenbacher, 2012; Raper, 2007).

GR is understood as a situational construct that emerges from the combination of a user's geographic information need (GIN) and the utility or influence of a geographic information object (GIO) (Raper, 2007). A GIO represents spatial features with a defined extent that may become relevant to a user. An example of a GIO is a POI.

Spatial and temporal attributes can be used for indexing and calculating the value of relevance to a GIO (De Sabbata & Reichenbacher, 2012; Raper, 2007). However, the process remains complicated, as GR contains both objective elements (e.g., spatial distance, opening hours) and subjective elements (e.g., user intent, preferences), making the GIN strongly dependent on context. As a consequence, a GIO is not uniformly relevant across users. This aspect makes it fundamental to know the map object properties and the vuser context (De Sabbata, 2013).

2.1.2 Understanding Context

Context has its origins in computer science and can be described as the set of conditions and attributes that shape the interaction between the system and the user. In the case of map usage, computer science refers to the varying set of conditions under which spatial information is accessed, interpreted, or acted upon (Abowd et al., 1999; Bartling et al., 2022). However, there is a universally agreed approach to modelling context (Bartling et al., 2023).

According to Bartling et al. (2023), context in mobile map applications can be systematically divided into three categories of extrinsic, intrinsic, and behavioural context. This taxonomy reflects an increasingly adopted structure in geographic information science (Fabrikant, 2023). Extrinsic context refers to external, observable factors such as location, time, movement speed, and the technological or environmental setup. Among these factors, location remains the most influential of these parameters in shaping user interaction and spatial decision-making (Sarjakoski & Nivala, 2005).

In contrast, the intrinsic factors capture internal user attributes, including user preferences or cognitive state. Recent studies have further highlighted the role of the individual in spatial abilities, such as the sense of direction and spatial anxiety, which shape the interaction patterns with adaptive spatial systems (Zingaro, Reichenbacher, Bartling, & Fabrikant, 2024). These factors are important for modelling the subjective dimensions of GR, as they influence how users perceive and prioritise spatial information (Fabrikant, 2023).

Last, the behavioural context includes physical movements and digital interactions, such as tapping, searching, or browsing. These behaviours can provide dynamic insight into user intent and tasks and provide information for real-time adaptation.

Capturing those attributes relies on both static and dynamic data sources, including GPS, mobile sensors, user interaction logs, and real-time data, such as weather conditions or traffic data, which can further enrich the modelling process (Bartling et al., 2023; Bartling et al., 2022; Bartling, Robinson, Resch, Eitzinger, & Atzmanstorfer, 2021).

While a wide range of context attributes can theoretically be considered, identifying and evaluating the most relevant ones remains a challenge, especially when accounting for the diversity of different user situations and goals in mobile mapping environments (Griffin et al., 2017). Understanding and considering these different dimensions is, therefore, a central part of context modelling and designing suitable map applications (Newman et al., 2010).

2.1.2.1 User Activities and Tasks

The contextual information not only characterises the surroundings of the map usage but also plays a significant role in shaping the activities of a user in the behavioural domain. Map interactions are mostly task-driven; understanding the intentions and reason behind this task is as important as knowing where and when a user is doing it (Reichenbacher, 2004).

Activity theory is a theoretical framework that helps describe user behaviour within a social and cultural context. t provides an important way to understand human activities as purposeful processes closely linked to their contexts. These activities are carried out through actions that transform objects or situations to achieve a desired outcome (Kaptelinin & Nardi, 2007). A user searching for food employs a sequence of intentional actions (e.g., checking options and navigating), which are embedded in a larger activity of satisfying a need. Here, GR becomes the system's skill in evaluating the GIOs that best support the user's goal.

With the available context information, GR operates as a dynamic system that aligns spatial information with the user's goal and activity, adapting a map accordingly. These goals can span a wide range of everyday activities, which are often tied to spatio-temporal cues. Empirical studies have shown that users frequently turn to maps to primarily answer practical, spatial questions and for receiving support in navigation, exploration, or task planning (van Elzakker & Ooms, 2017). Zingaro and Reichenbacher (2022a) have found that mobile map usage frequency increases with distance from home, which highlights the relevance of maps as orientation tools in unfamiliar environments. However, real-world usage can differ greatly, as users often shift between exploratory and goal-directed behaviour, depending on situational needs and device affordances (Savino et al., 2021).

Modelling such activities, however, is complex and often difficult in a straightforward way. To address this challenge, previous studies have examined how extrinsic factors, such as the time of day and day of the week, play a significant role in user behaviour. Their findings provide valuable insights for context modelling and help explain where users go and what they do. For example, research has shown that user activities are structured around daily and weekly routines: food-related activities happen mostly during lunchtime and shift later at weekends (Teixeira, Da Cunha, Azeredo, Rinaldi, & Crispim, 2024). Shopping, on the other hand, varies with the day of the week and is more popular on the weekends (Hirsh, Prashker, & Ben-Akiva, 1986). Commuting patterns are also dependent on the day of the week and have two high peaks in the mornings and evenings (Hatcher & Mahmassani, 1992; Li, Guensler, Ogle, & Wang, 2004). Leisure and cultural activities tend to happen mostly on weekends and/or in the evenings (García, 2017). Finally, outdoor engagement peaks in the early afternoon, coinciding with daylight availability, and occurs more frequently on weekends (Haswell, West, Savill, Whitham, & Martin, 2014; Hayir-Kanat & Breuste, 2020). These patterns can directly inform the GR calculation and adaptive prioritisation of POIs.

2.1.2.2 Spatio-Temporal Mobility

User mobility is another major contributor to modelling the user context and can be split into spatiotemporal proximity and directionality (De Sabbata & Reichenbacher, 2012; Huang, Ma, & Liu, 2015). Spatiotemporal proximity captures how well a geographic object aligns with the user position, destination, and time constraints. An ideal object would lie on the user's route and fit in the available timeframe, making the combination of the spatial and temporal proximity a reversed measure of utility (De Sabbata, 2013; Reichenbacher, De Sabbata, Purves, & Fabrikant, 2016). Building on this idea, research has also aimed to address the limit of the travel distance, finding significant variability according to the mode of transportation. By walking, those limits can range from 0.5 to 2km and go beyond 20km when using a car (Papadimitriou, Theofilatos, Yannis, Sardi, & Freeman, 2012; Shu et al., 2024). Factors such as mode of transport, time of day, or movement speed further influence the mobility component (Elias, Hampe, & Sester, 2005). For example, a faster mode of transportation allows reaching more distant GIOs in the same time window.

Directionality, in turn, refers to how well a geographic object aligns with the user's intended travel direction. Objects located along the route require fewer detours and therefore are perceived as more relevant (De Sabbata & Reichenbacher, 2012). This is closely tied to time geography, which studies the interaction of time, space, and user activity, as well as the resulting constraints or enabling of user behaviour (Miller, 2004). User activities require a movement through space and time and are limited by fixed locations and time budgets (Miller, Wu, & Miller, 2001). The example of the activity 'getting food' only succeeds as long as the restaurant is reachable and open on time.

2.1.2.3 Geographical and Physical Environment

Physical environmental factors, such as weather, user surroundings, or ambient light, also play a role in shaping the user context (Sarjakoski & Nivala, 2005; Yu, Zu, Wang, Yu, & Li, 2019). Unfavourable conditions, such as bad weather or darkness, can impact a user's visibility and influence their comfort level and influence users' willingness to travel further distances. Those factors form an objective, extrinsic layer of the user context.

Beyond those atmospheric elements, the spatial arrangement of GIOs also contributes to the environmental context. This factor is defined as the geographic environment and considers, how the entities are spatially arranged within their surroundings (De Sabbata, 2013). A restaurant, therefore, might be more relevant either by having a lot of other restaurants near to each other or by offering alternatives or by having complementary entities located closely (such as hotels).

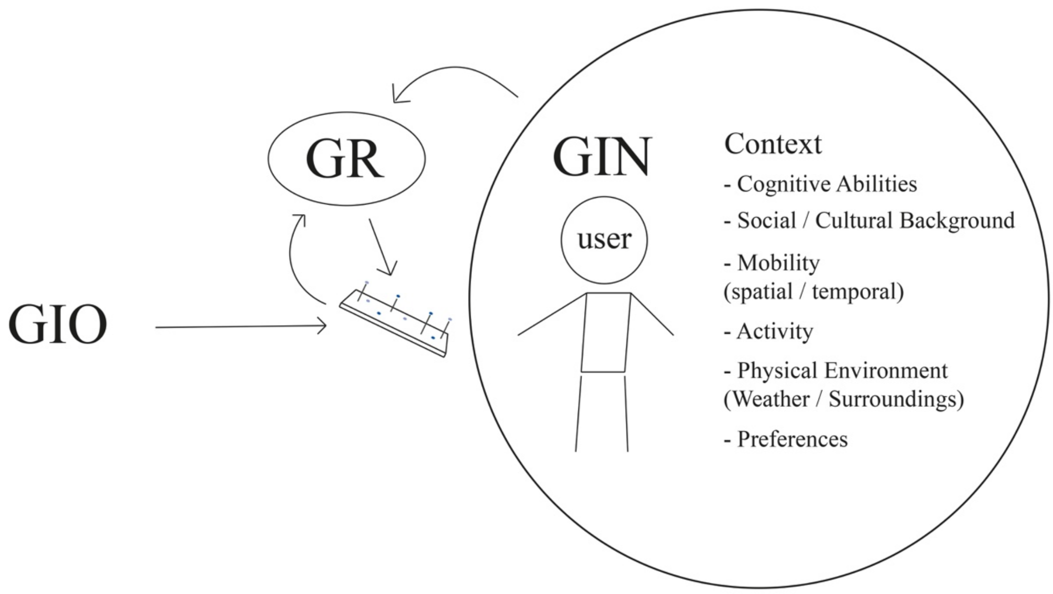

Figure 1: Visual representation of the factors of geographic relevance and user context

2.1.2.4 User Preferences and Background

User preferences are another factor of user context that directly influence the perception of the relevance of GIOs in a given situation (Bartling et al., 2022). It reflects the domain of the intrinsic context and relates both to the user activity and the thematic or topical interests associated with it (Rashid et al., 2002; Yin et al., 2015). Beyond the situational preferences, users have spatial anchors, which represent the structure of their mental map of geographic space (Couclelis & Tobler, 1987). These places, such as the workplace, home, or frequently visited areas, shape the movement patterns and define zones of relevance. Figure 1 provides an overview of the addressed subset of context variables and their conceptual relationship with GIN, GR and GIOs.

2.1.3 Assessing Geographic Relevance

Based on the highlighted contextual variables, the user situation can be modelled, allowing the assessment of the GR of individual GIOs. The GR functions as a computational bridge between user needs and spatial data, enabling a dynamic evaluation of which GIOs best support the user's current goal and behaviour. The resulting relevance scores are thus derived from a computational function that considers both the attributes of GIOs and the user's context:

$$GR\ (\ GIO, Context) \rightarrow Relevance\ Scores$$

The structured conceptual model of the GR assessment proposed by De Sabbata (2013) bridges theory and implementation. By integrating the different dimensions of user interaction into a coherent method for calculating GR, it provides a strong foundation for implementation into an AMMA. Earlier assessments of GR modelling often relied on heuristic rules, spatial filtering, or static thresholds (Ballatore, McArdle, Kelly, & Bertolotto, 2010). The interest scores of a GIO were calculated on its spatial proximity to the user, combined with the user's interaction history. The latter is modelled through click frequency and adjusted by a time-decay function to account for the loss of relevance over time.

The De Sabbata model, however, offers a more nuanced, multi-dimensional representation of user context and demonstrates the interaction and influence of GR, GIN, and GOIs. The basis of the chosen factors was evaluated empirically under the assumption that a semantic distance value can be calculated for each criterion by comparing the GIN of the user to the GIO. These values can then be combined into a normalised score [0,1] that serves as an estimator for GR. This approach is in line with the first law of geography (Tobler, 1970, p. 236), which states that 'everything is related to everything else, but near things are more related than distant'. In this model, the geographic space extends beyond the physical to include semantic, contextual, and temporal domains. Accordingly, entities with a shorter semantic distance are considered more relevant (De Sabbata, 2013).

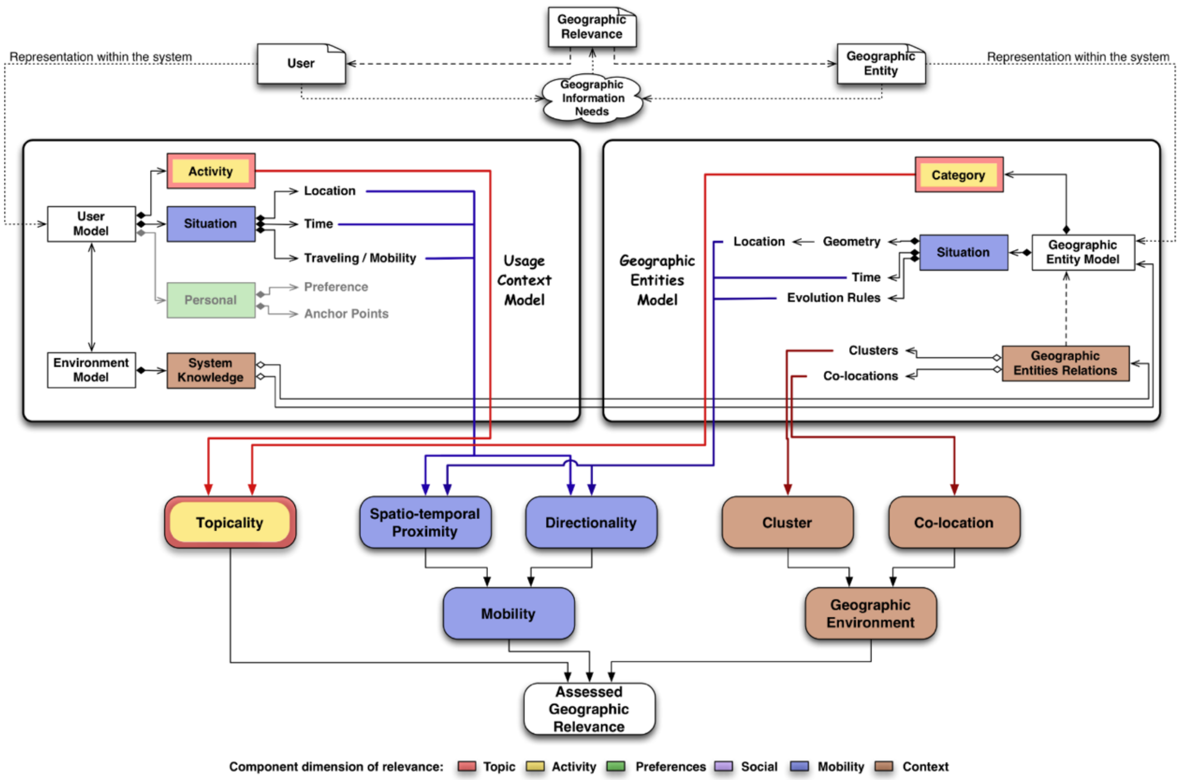

Figure 2: Computational framework for the assessment of geographic relevance (De Sabbata, 2013)

The proposed model of De Sabbata (2013) computes GR based on five key elements, each reflecting the distinct dimension of the user context. Topicality should reflect how well the semantic category of the entity matches the user's current information need or activity. The dimension tries to account for the topical and activity components of the user context. The spatio-temporal proximity and directionality are used in combination to assess situational elements and model user mobility (De Sabbata, 2013). Geographic environments are considered in terms of clustering and co-locations of the GIOs. Both factors were shown in crowdsourcing to be central to context modelling and assessing GR (De Sabbata & Reichenbacher, 2012). Clustering serves as an indicator if the entity is located near other, similar entities (e.g., several other restaurants besides the restaurant itself). It is assessed by the size of the cluster and the distance of an entity to another entity of the same amenity. The first reflects the aspect that larger clusters suggest better availability of similar options. The second emphasises that users prefer dense groupings of similar entities, with closer distances rated higher (De Sabbata, 2013). Co-location, on the other hand, considers complementary entities in the surrounding environment. A co-location rule links categories of premises entities to a category of conclusion entities, indicating that entities of the former are often found near the latter. Co-location is evaluated using two main components: the count of nearby entities from the conclusion category found within a distance threshold from the entity in question, and the proximity to the nearest such entity. Many complementary features nearby indicate a supportive environment for the user's activity, resulting in a higher co-locations score (De Sabbata, 2013).

The resulting three main dimensions, as visualised in Figure 2, can then be combined to achieve an assessed GR. De Sabbata (2013) has suggested a continuous preference logic (CPL) to weigh the differing importance of the individual criteria. CPL offers flexible operators to blend in mandatory and desired inputs. Within the mobility component, spatio-temporal proximity is considered mandatory, while directionality is treated as desired. This score is combined with the mandatory factor of topicality and the desired geographic environment component. Together, these calculations form a critical foundation for building an adaptive mobile map application, where the map's content can be filtered and prioritised based on the user's real-time context.

2.2 Mobile Map Application Design

A practical design of mobile map applications requires a careful integration of technological, cognitive, and visual principles to ensure usability and accessibility. This section builds on the foundations of GR and presents a framework for adaptive mobile map design.

HCI forms the theoretical and practical foundation for the design of user-centred systems. The focus is on human perception, processing, and response to digital interfaces (Beaudouin-Lafon, 1993). In mobile cartography, HCI provides tools and a framework for designing context-aware, visually apparent, and cognitively manageable interfaces that adapt to changing user tasks and preferences, as well as supporting dynamic interaction models that offer relevant information without cognitive overload (Reichenbacher, 2004).

The section begins with the HCI design principles, which set the basis for user interface (UI) design and information flow. Various design principles are introduced to establish the theories of perceptual structure and visual clarity. Finally, adaptive design principles are discussed that can support different users and use cases.

2.2.1 Design Principles

An application design needs to be useful. This usefulness can be assessed by the combination of its utility and usability. The first refers to the extent to which the necessary functions and tools are provided to meet the goals of the user activity. Usability, on the other hand, concerns the interaction quality with the tools provided and describes how intuitive, efficient, or satisfying the experience is (Ayada & Ezz Eldin Hammad, 2023).

A helpful design must support the formation of an accurate mental model in the user through consistent and meaningful interactions. This ensures that the user understands how to operate the interface and can do so with minimal friction. Such an alignment between system behaviour and user expectations is the basis for an effective engagement in real-world contexts.

2.2.1.1 User Discoverability

A user should be able to identify all possible interactions within an application and understand how to perform them without requiring any prior knowledge. This quality, known as discoverability, is essential in supporting both usability and the development of an accurate mental model of the system. As Mackamul, Casiez, and Malacria (2024) describe, discoverability is not a single feature but an essential property of the system and its interface design. The design should include intuitive cues and structures to facilitate user learning based on these design principles. Affordances, for example, suggest potential use cases of objects to the user (Norman, 2013). When not effectively leveraged, users may overlook key functionalities (Wingrave, Bowman, & Ramakrishnan, 2001).

Signifiers can help with communication to the user. These are perceivable indicators that guide user behaviour, such as drop shadows, button shapes, labels, or colours. An appropriate symbol can also be a signifier of an affordance (Ayada & Ezz Eldin Hammad, 2023).

The integration of constraints helps guide user behaviour through the regulation of specific actions within meaningful boundaries. With that, possible errors can be prevented by relying on commonsense actions (Norman, 2013). Providing clarity and visibility to the functionalities makes them more exposed to the user and increases their noticeability (Mackamul et al., 2024).

During or after interaction with a feature, feedback mechanisms, such as error messages or loading signs, are vital tools to keep users informed about the system status and reduce uncertainty. The responsiveness of an interface is important, since long loading times or errors lead to user frustration (Ayada & Ezz Eldin Hammad, 2023; Mackamul et al., 2024). HCI findings even have shown that a higher responsiveness of applications correlated with higher task success rates and lower completion times (Caro-Alvaro, García-López, García-Cabot, & Mavri, 2025). Attig, Rauh, Franke, and Krems (2017) have summarised and pointed out benchmark numbers that should be fulfilled for a responsive mobile application. A common upper limit for 'instantaneous' feedback is a latency up to 100ms. An initial load should happen with a maximum latency of one second. If latencies exceed 2s, users expect visual loading indicators. Seow (2008) found, that users will lose focus after more than 10s for continuous interactions such as system computations.

Furthermore, consistency among the different interface designs is a priority. A uniform use of fonts, colours, icons, and interaction patterns across an application reinforces user expectations, supports learning, and helps build the user's conceptual model (Ayada & Ezz Eldin Hammad, 2023). By aligning with conventions, prior knowledge of similar systems can be reused, and the interaction can be facilitated (Mackamul et al., 2024).

A system should appear as simple as possible with minimal visualisation clutter (Tufte & Graves-Morris, 1983). In a mobile context, unnecessary or irrelevant information should be minimal. This also applies to geovisualisation, where filtering and hiding irrelevant features, generalisation, or aggregation can help reduce complexity (Crease & Reichenbacher, 2011). A clear visual hierarchy can structure the displayed content into the background (e.g., basemaps) and foreground (POIs). The centre–surround principle suggests that users view the background broadly and the foreground in more detail, offering a source of inspiration. This helps with navigation and focuses perception to improve search efficiency and orientation (Crease & Reichenbacher, 2011; Swienty, Reichenbacher, Reppermund, & Zihl, 2008). The spatial proximity of features within this hierarchy also influences perception: closely grouped features and more densely populated maps improve user comfort and increase success rates in finding relevant information (Bartling et al., 2021).

One practical application of these principles is the use of visual variables to assess GR (see Section 2.2.1.3). Another is modality, where the visual hierarchy is reinforced through separate, focused views. Modal presentations allow users to switch from a general overview to an isolated interaction with selected content, reducing clutter while supporting targeted exploration.[1]

[1] https://developer.apple.com/documentation/swiftui/modal-presentations, accessed 25.09.2025.

2.2.1.2 Gestalt Laws for Interfaces

Designers can apply the Gestalt laws of perception to organise application interfaces in a way that allows humans to make sense of complex data. These interfaces can be visually pleasing and support discoverability (Norman, 2013). Figure 3 presents a selected overview of those principles.

Symmetry is one of the Gestalt laws and refers to the human tendency of perceiving symmetrical arrangements as unified and aesthetically pleasing. In interfaces, symmetry often appears in grid views or evenly spaced lists, generating visual balance through a predictable layout. When elements are repeated consistently, this is designated as translational symmetry. This type of symmetry can be combined with intentional asymmetry within the symmetrical structure to strategically draw attention to the centre of the interface and the main content without breaking the overall coherence (Butler, 2023a).

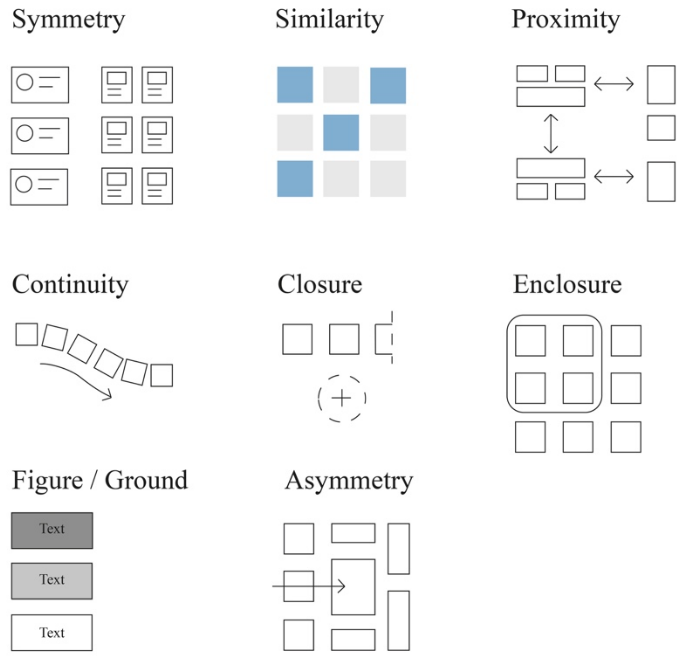

Figure 3: Overview of the gestalt laws for interface design

Similarity explains how elements that share the same characteristics (colour, shape, size, orientation) are perceived as part of the same group, even when separated through space (Graham, 2008). This design supports cohesion and helps users find content categories or related functionality. The **proximity** of elements to each other can also serve as a grouping indicator. Items that have a position close to each other are recognised as the same group, while those further away from each other are perceived as unrelated. When elements are bound with the same area or boundary by **enclosure**, they are again perceived as a group belonging together. This can be achieved through background colours, shading or lines (Rock & Palmer, 1990).
The Gestalt principle of **continuity** describes the human tendency to follow continuous paths through alignments of elements across the screen. It creates a sense of connectedness and

coherence, especially in sequential or content-heavy layouts (Chang, Dooley, & Tuovinen, 2002; Graham, 2008). Incomplete shapes or patterns are filled in with the missing information by the human brain. This **closure** allows users to recognise familiar forms when parts of them are obscured and could be used to, for example, signal the availability of more information (Butler, 2023b; Chang et al., 2002; Graham, 2008). **Figure/ground** helps the brain distinguish an item (figure) from the underlying background (ground). Using contrast can ensure the figure elements (text, button, links) clearly stand out (Graham, 2008).

2.2.1.3 Visualising Geographic Relevance

While the Gestalt principles shape the overall interface design, the cartographic design is responsible for communicating the assessed GR of GIOs to the user. With the help of the principle of conciseness (see Section 2.2.1.1), visual variables, and the understanding of graphical perception, the POI attribute of relevance can be encoded and represented visually. These variables were initially introduced by Bertin (1983) and have been further refined over time. Each variable serves a unique perceptual function and supports ordering, differentiation, or quantification of features. The visual information is processed automatically and segments visual stimuli into meaningful patterns even before conscious attention is applied. This pre-attentive processing allows the user to detect salient features immediately without deliberate effort. Map symbols that leverage pre-attentive features with visual variables are more likely to stand out and be grouped intuitively (John, 2014).
The most obvious and simple of those variables is the geographical location of the entity. However, to communicate a deeper meaning (such as GR), the variables must also be able to communicate in an ordinal (or even better, quantitative) characteristic. Michael Förster (2023) has analysed the visual variables in a mobile context and found that colour (value and saturation), transparency, and size offer the needed resources to make a feature recognised as relevant by a user. Those four variables are visualised in Figure 4.

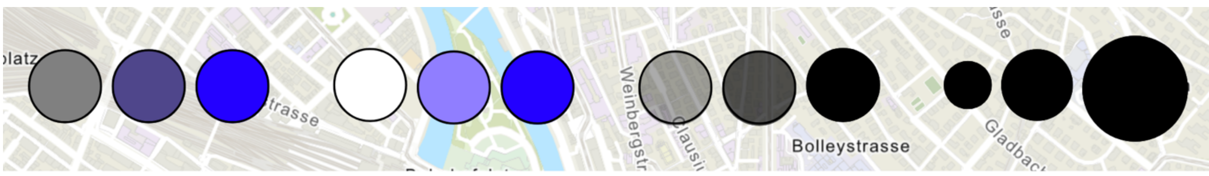

Saturation Value Transparency Size

Figure 4: Overview of the ideal visual variables for communicating relevance

Colour value describes how light or dark an element appears, while the saturation reflects how much grey is mixed into a colour. These two variables can be detected quickly and effectively (Cláudio, Carmo, Labmag, & Mendes Leal, 2011; Garlandini & Fabrikant, 2009). However, size and transparency were found to perform even better in communicating relevance.
Transparency refers to the degree of visibility of the background through the symbol. Research has it to be particularly effective: Crease (2014) demonstrated that opacity is well suited for conveying GR intuitively, while Oliveri and Reichenbacher (2021) found that transparency enables users to identify and order POIs more accurately and rapidly than with colour hue and value. Size, in turn, is defined as the amount of space a label occupies on the screen (Roth, 2017). This is a powerful variable due to its capability of allowing for quantitative interpretations by users. Larger labels naturally draw more attention and were found to be the most effective in relevance assignments (Cláudio et al., 2011; Garlandini & Fabrikant, 2009).

2.2.1.4 Accessibility

Designing with accessibility in mind is both an ethical obligation and a practical necessity. Users interact with the interface in various ways. They differ in cognitive, sensory, and physical abilities. Effective interface design must account for these differences and ensure that all users can use the application. A primary consideration should be the cognitive processing capability, which varies significantly across users (Steichen & Fu, 2019). Atzmanstorfer et al. (2016) and Bartling, Resch, Eitzinger, and Zurita-Arthos (2019) have observed that users with limited experience in smartphone and map application usage struggle with orientation tasks. These difficulties arise from the need to process dynamic visual information while maintaining spatial awareness and aligning map content with real-world contexts. To support such users, interfaces and map designs should therefore be minimised in complexity, supporting the design principle of simplicity and discoverability.
The colour encoding of the interface is another critical factor for users with colour blindness. To ensure accessibility, interfaces should maintain high colour contrast to help users distinguish elements. Furthermore, colour combinations that are commonly indistinguishable to individuals with colour blindness should be avoided (Ayada & Ezz Eldin Hammad, 2023). A common example for 8% of the male population is the difficulty to distinguish between red and green colours (Jenny & Kelso, 2007).
Beyond visual limitations, colour choices carry psychological and cultural connotations. Red is often perceived as negative or subordinate, while green is associated with acceptance or

positivity. Those meanings can unintentionally influence user perception and should be addressed. However, emotions associated with colours are not universal and vary across different cultural backgrounds. Tools such as Colour Brewer can serve as a good guide on colour combinations that are both accessible and perceptually distinct (Coalter, 2020). A good practice is to use labels and symbols as an addition to the element of colour to ensure the interpretability of information.

Not only the colour but also the readability of text elements plays a fundamental role in accessibility. Texts should be large enough with appropriate spacing to read comfortably, and legible fonts should be chosen (Ayada & Ezz Eldin Hammad, 2023). This can support users with visual impairments and enhance readability on screen under varying environmental conditions, including fluctuating luminance and screen reflections.

For users with hearing impairments, the support of auditory information with visual cues or captions is essential. For users with motor impairments or those operating the application in constrained environments, multimodal interaction (gesture and voice input) enhances inclusivity and usability (Packia et al., 2024) .

An inclusive design improves the usability of applications for people with disabilities and benefits all users by providing a clear, flexible, and straightforward application interaction. The cognitive load can be minimised, and the interaction between the application and user can happen in a way that respects capabilities and preferences.

2.2.2 Adaptivity

Adaptive map design enables the tailoring of the adaptive map interface to different users and their specific contexts. It indicates the system's ability to automatically adjust its behaviour in response to user preferences, behaviours, or environmental conditions. It can happen either preemptively or in real time during usage and is set automatically by the system itself.

The concept of adaptive systems aims to reduce cognitive load and considers the most constrained-use scenario. On the small screens of mobile map applications or with limited user attention, adaptive behaviour is vital for fast and easy understanding of geographic information (Roth, 2019). There are three main types of adaptation for map applications: the geographic information that is represented, how this information is visualised, and the application interface (Reichenbacher & Bartling, 2023).

These forms of adaptation become particularly important in dynamic usage contexts. For example, when movement speed increases, decision-making must happen quickly, which emphasises the need for adaptive visualisation with information tailored to immediate requirements (Meng, 2005; Reichenbacher, 2005). Similarly, the adaptation of colour schemes in low-light conditions can improve readability and reduce eye strain in dark environments is another example of an adaptive interface (Zhang, Adipat, & Mowafi, 2009).

In comparison to static map designs, adaptive interfaces improve task performance, usability, and user satisfaction (Packia et al., 2024). Personalisation further strengthens this effect, as maps can learn from usage patterns (such as frequently visited places or saved favourites) and prioritise more relevant information by assessing GR (Ayada & Ezz Eldin Hammad, 2023). Features such as automatic aggregation of map content based on zoom levels and GR help maintain visual clarity by adjusting the level of detail presented to the user (Reichenbacher, 2004). In this sense, cartographic generalisation can be reimagined in a context-dependent way, allowing the map to adapt dynamically to user needs (Roth, 2019).

Building on the role of adaptation and personalisation, GR provides a key mechanism for tailoring map interfaces to the user. By evaluating GR, both the interface and the map content can be adapted to align with the users' immediate needs and context, maximising usability on the constrained screen space of a mobile device. Including a diverse set of recommended places and not only sticking to certain categories further enhances usability (Chen, Yu, Tang, He, & Zeng, 2017; Liu, Fu, Yao, & Xiong, 2013). An optimal approach may, therefore, involve offering a balanced mix of high-priority, contextually relevant locations and landmarks to enhance spatial orientation and navigational support.

However, timing and transparency are crucial since automation always comes with a loss of control for the user. Adaptive behaviours should be event driven and context aware, rather than constant or arbitrary. Crease (2014) showed that adaptations triggered too early risk confusing users, while those triggered too late lose their effectiveness. To be successful, an adaptation should happen in a way that keeps the user in control, fostering trust and maintaining situational awareness (Bartling et al., 2022; Crease, 2014). A lack of transparency or the user having a sense of being 'overridden' by the system affect the user experience negatively (Kiefer, Giannopoulos, Athanasios Anagnostopoulos, Schöning, & Raubal, 2017). Adaptive visualisation should therefore balance between intelligent automation and user regulation, ensuring that the users understand why and how the map application changes. With those design principles in mind, the UI and map content itself can be designed and displayed in a way that prevents misinterpretations, reduces cognitive load, and highlights the most relevant GIOs to the user in their current context.

2.3 Technical Solutions for Mobile Maps

2.3.1 Mobile Map System Architecture

Mobile devices are now the main entry point for a variety of services. Despite their widespread use, they have limitations in processing power, screen space, and battery life, making computationally expensive tasks difficult to execute locally. To overcome these constraints to some extent, mobile map systems are increasingly relying on cloud computing architectures.

Cloud computing facilitates offloading data procession tasks to remote servers, allowing applications to perform real-time computations and deliver context-aware content on demand. In particular, software as a service models allow users to access map-based applications without managing the underlying infrastructure. This can enhance the responsiveness and enable smooth data updates as well as personalised spatial content delivery (Xhafa, 2018).

Most AMMAs follow a client-server architecture, where the mobile client collects contextual data (e.g., location, user preferences) and communicates with the server to retrieve tailored content. This dynamic data exchange enables personalised map representations and context-aware responses from the server.

In a real-world observational study, Savino et al. (2021) demonstrated how such client-server architecture functions during a diverse set of usage scenarios. While the MapRecorder collected the interaction data of the user, the application was used for often quick and fragmented sessions. These findings support the need for a responsive architecture to adapt quickly to the rapid shifts in user attention and behaviour.

Despite their benefits, cloud-based systems rely heavily on a stable network connection. Interruptions can lead to latency or reduced functionality, which can be particularly problematic in outdoor or remote environments. Caching and batch-loading techniques can be applied to reduce the need for continuous data requests. This reduces the number of client-server requests and improves the application speed by minimising data retrieval and energy consumption (Naeem, Yahui, & Abrar, 2025). Bottlenecks can be bypassed with the caching of frequently needed content (Durand, Ma, Pinnecke, & Saake, 2018).

2.3.2 Object-Oriented Design and Software Patterns

Developing complex adaptive systems for mobile platform necessitates the use of maintainable and extensible software paradigms. Object-oriented programming (OOP) plays a critical role in this regard by enabling modularity, encapsulation, and reusability, all of which are essential for separating concerns across data layers, UI rendering, and behavioural logic (Beck & Cunningham, 1989; Onu, Ikedilo, Nwoke, & Okafor, 2015).

The principle of OOP is based on the object, a software entity that encapsulates data and functions. The program then consists of objects that interact with one another (Lewis & Loftus, 2018). In AMMAs, OOP allows developers to encapsulate components, such as context sensing, GR scoring, map rendering, and UI adaptation, into discrete modules. This encapsulation makes it easier to test and maintain the individual systems components, such as switching a GR model or adapting the visualisation layer, without impacting the system core.

OOP provides proven design patterns to help design a functional system. The specific design patterns implemented in the system are introduced in Section 3.5.2. They include foundational patterns, such as the Observer (for real-time context updates), Strategy (to switch GR calculation methods), and Model-View-View-Model for clean UI separation. These design choices enable real-time interaction handling, maintain interface state coherence, and support reactive updates as a key requirement for AMMAs (Miiggel, Rhol, Speicherl, Bihlerl, & Cremersl, 2007).

Previous implementations, such as RecoMap or A-POInter, have shown how modular and pattern-based software design can enhance adaptive map applications. These systems rely on layered logic to separate data retrieval, application logic, and UI presentation (Ballatore et al., 2010; Hill & Wesson, 2010), making them suited for iterative prototyping and feature extensions.

2.3.3 Data Privacy

Adaptive systems are inherently data-driven, relying on extensive user context and tailored content. However, the benefits of personalisation must be weighed against the increasing public concern about data privacy. Users are often hesitant to share behavioural data, particularly when they feel a lack of control over how their data is used or stored (Awasthi, Tewari, & Mallick, 2025; Ul Haque, Khan, & Fahim, 2023). Recent research has suggested that trust in service providers is strongly tied to the availability of privacy controls over data usage (Ul Haque et al., 2023). The distinction between sharing (who can see the data) and usage (how the data is exploited) is often fuzzy, leading users to overestimate their data protection.

To mitigate these concerns, the use of on-device processing and local storage models is becoming best practice. Context sensing and relevance calculations occur locally, keeping personal data confined to the device and minimising unnecessary data exposure. This model can reduce privacy risks and boost responsiveness in GR computation.

Nevertheless, trade-offs remain between usability, automation, control, and privacy (Bartling et al., 2022). Adaptive systems must strike a balance between demands to avoid undermining user trust and requiring data logging or context tracking for adaptation mechanisms and evaluation purposes. Designing for transparency enables users to understand how and why their data affects the interface.

3 Methodology

This chapter presents the design process and application architecture of the mobile map application, which is designed to support adaptive user experiences, based on the concept of GR. Building upon the theoretical foundations outlined in Chapter 2, the methodological objective is to translate these abstract ideas and models into a concrete, usable prototype. The goal is to demonstrate the technical feasibility of real-time context-aware adaptation and to establish a modular and extensible system that can serve as a foundation for future research and scientific experimentation.

3.1 Design Goal and Use Cases

The primary objective of the prototype was to support mobile users with spatial information that is contextually relevant to their current activity and real-world situation. The system was built upon the assumption that users often operate in dynamic, unfamiliar environments, where mobile devices serve as a critical tool for planning and decision-making. These scenarios often happen with external distractions and under time constraints.

The design concept was rooted in personalisation and contextual awareness, following the theoretical principles introduced in Chapter 2. To guide both the logic of the application and its adaptive functionality, three central use cases were defined. These were based on the findings that users typically turn to maps to answer questions related to navigation, exploration, and planning (van Elzakker & Ooms, 2017). The selected use cases reflected these fundamental map-use goals and corresponded to the following distinct cognitive tasks:

- **Use Case 1:** On-the-Go Relevance Discovery: This scenario addressed real-time exploratory behaviour, where a user visits an unfamiliar location and seeks information about relevant nearby places, such as restaurants or shops. It aligned with the common map task of answering ‘What is nearby?’, which requires the system to compute and display contextually fitting POIs that match the user’s immediate needs and location
- **Use Case 2**: Revisiting and Consulting Favourite Places: Previously saved locations of places can be used by the user to retrieve specific information, such as checking opening hours or other POI attributes. This reflected reference-based tasks where the map serves as a memory aid, helping users reorient or carry out follow-up actions with minimal friction.

- **Use Case 3:** Location Exploration and Planning: When preparing for a future visit, users can search for and explore unfamiliar areas in advance. The system supported this behaviour by offering a search function, automated semantic filtering, and delivering context-sensitive POI recommendations.

These use cases were chosen to span immediate, reflective, and prospective modes of spatial reasoning, ensuring that the prototype could adapt to a wide range of real-world interactions.

3.2 Requirements and Scope

While the previous section established the conceptual goals and intended use cases of the system, this section outlines the technical and scientific requirements necessary for its implementation. These requirements fulfil the application's dual role as both a functional mobile tool and a platform for scientific research and evaluation of adaptiveness in mobile cartography.

3.2.1 Functional and Technical Requirements

Based on the theoretical groundwork laid in Chapter 2 and the practical use cases described in Section 3.1, the functional and technical requirements of the AMMA MyPlaces were defined. These requirements translate the conceptual foundations into actionable system specifications, ensuring that the application design remains both theoretically grounded and practically applicable. They serve as the basis for guiding implementation choices and for evaluating whether the system fulfils its intended purpose. The functional requirements derived are shown in Table 1.

Table 1: Overview of the functional requirements for the adaptive mobile map application prototype

Code	***Requirement***	***Description***
F1	Context awareness	Sense and process spatial, temporal and behavioural data for context-aware decision making
F2	Activity modelling	Assume user activity with extrinsic variables to then use for GR computation
F3	Relevance computation	Implement a GR computation model that uses variables from topicality, mobility and geographic environment

F4	Adaptive visualisation	Dynamically adjust POI visibility and basemap layers based on the computed context
F5	Visual feedback & overrides	Provide feedback on adaptations and allow manual override mechanisms to support user control
F6	User profiles	Support multiple user profiles with individually stored favourites and interaction history
F7	Search	Implement search functionality to look up places by address names
F8	POI popups	Allow for further IR of POIs through popups with the feature attributes
F9	Orientation	Provide aid for user orientation with the display of the current user position and surrounding landmarks on the basemap

The technical design must ensure responsiveness, scalability, and extensibility across a variety of user conditions. To achieve this, the technical requirements specify the underlying capabilities needed to support the functional aspects of the system and to guarantee reliable performance in real-world use. The core technical requirements are shown in Table 2.

Table 2: Overview of the technical requirements for the adaptive mobile map application prototype

Code	***Requirement***	***Description***
T1	Object Oriented Design	Use OOP principles and design patterns for modularity and extensibility
T2	Efficient Data Management	Employ local storage for caching and real-time relevance computation
T3	Performance Optimisation	Ensure quick responses for GR scoring, interface updates and sensor integration
T4	HCI Design	Follow HCI best practices to support the application discoverability
T5	Privacy	Store data locally and don’t expose any sensitive data to the server

3.2.2 Scientific Research Requirements

Beyond a role as a functional tool, the application is also a platform for scientific inquiry into context modelling, GR assessments, and adaptive map design. Its implementation is driven by the need to address key research questions in cartographic adaptation and mobile user interaction. The application must first act as a demonstrator for theories of GR, particularly the model proposed by De Sabbata (2013). This includes integration of user activities and adapting the interface based on changes of context. These activities allow the testing of GR models in real-world conditions. A crucial requirement is the evaluation of context sensing. The system must go beyond static extrinsic factors and explore more nuanced variables, such as behavioural intent, preferences, and dynamic goal shifts. Bartling et al. (2023) have highlighted the lack of tools that integrate such complex inputs in GR assessment. Another major research requirement was the evaluation of adaptivity itself. While Crease (2014) has shown that adaptive map applications can be more usable than static map representations, real-word field studies are needed. The scientific requirements are summarised in Table 3.

Table 3: Overview of the scientific requirements for the adaptive mobile map application prototype

Code	***Requirement***	***Description***
S1	Experimental Validity	Log GR scores, user context, and behaviour to support analysis
S2	Reproducibility	Version Control and clear implementation documentation
S3	Evaluation Metrics	Enable export of interaction data for further analysis
S4	Modularity	Ensure implementation structure for further future extensions and updates through OOP

By meeting these research requirements, the application aims to become a fully featured experimental tool for advancing the field of adaptive cartography.

3.2.3 Application Scope

To maintain a focused and manageable implementation scope, the application was geographically constrained to the national extent of Switzerland. This limitation ensured sufficient spatial and environmental variability for testing adaptation logic across urban, suburban, and rural contexts, while keeping data hosting and processing within feasible bounds. The setting also aligns with the research environment of the GIVA group and ensures real-world relevance for testing.

The application was designed as a proof-of-concept implementation, supporting only a select set of thematic user activities. As discussed in the theoretical background of this thesis, common daily behaviours, such as dining, shopping, commuting, leisure, and outdoor activities, are influenced by temporal and behavioural context variables (García, 2017; Hirsh et al., 1986; Teixeira et al., 2024). Rather than attempting to model all possible user intents, the prototype focused on a curated subset of these introduced thematic activities to ensure depth of implementation and experimental testability. These activities were selected based on their frequency, potential for modelling, and relevance to the use cases outlined in Section 3.1.

The prototype's level of adaptivity also reflects its exploratory role. It is not a complete commercial system but rather a research-grade foundation designed to demonstrate and test the feasibility of adaptive behaviour based on contextual inputs. Social and cultural patterns as well as cognitive abilities are, therefore, not yet considered in the adaptation process. However, the system's modular architecture allows for future extensions and refinements of relevance models (RMs) as the framework evolves.

Technically, the implementation targeted the Apple iOS platform, specifically iPhones, due to advantages in both practical and technical considerations. iOS offers predictable performance and consistent hardware specifications, which makes the development and performance optimisation easier. Compared to the second most widely used operating system for mobile applications in Switzerland, iOS features significantly less hardware and OS version fragmentation (Sun, Chen, Liu, Grundy, & Li, 2023; Xu, Chen, & Nie, 2014), reducing the risk of compatibility issues and ensuring a more consistent user experience across devices. Studies have shown that native iOS applications tend to be more energy-efficient while offering performance that is comparable to their native Android counterparts (Gyorödi, Zamaranda, Georgian, & Gyorödi, 2017). With a large user base and clear documentation, iOS offered a stable and efficient foundation for building an AMMA prototype.

3.3 Context Modelling Strategy

The implementation of GR requires sophisticated context modelling that captures the multifaced nature of user-environment interaction. Building upon the computational framework of De Sabbata (2013), this section presents the modelling strategy that translates the theoretical concept into operational algorithms capable of real-time relevance assessments.
The main challenge was the inference of the user intent without explicit input. The system, therefore, had to determine the user activity and interests solely based on the observable context. This required modelling both the user's thematic interest (what type of places they seek) and the relevance of individual places, given those interests. To address this, a dual machine learning (ML) model architecture was developed:

- Thematic model (TM): predicts activity-based interests from context
- RM: scores POIs based in multiple relevance dimensions.

The strategy pipeline began with the retrieval of the raw data. Afterwards, rule-based scoring functions for both the relevance and thematic scores were implemented. Based on those scores, a synthetic dataset was generated, which then served as a ground truth to train the two final GR assessment models. That way, model training could easily be repeated, improved, and implemented in the final application. All the modelling was implemented using Python.

3.3.1 Context Modelling Data Sources

The foundation of context modelling relies on three primary data sources that provide geographic, environmental and temporal information.
Points of Interest Data: The POI dataset derives from OpenStreetMap (OSM) data for Switzerland and was downloaded from Geofabrik.[2] The raw OSM data was pre-processed and controlled using ArcGIS Pro to then be directly published as a hosted feature layer on ArcGIS Location Platform (Fig. 5). Therefore, efficient spatial queries via application programming interface (API) calls and automatic data updates, without application modifications, could be made.
While OSM data provides comprehensive coverage, its open-source nature introduces uncontrolled variability in data quality and completeness. Rural and outdoor areas have a lower

[2] https://download.geofabrik.de/europe/switzerland.html, accessed 22.01.2025.

POI density and less consistent attribute completeness, particularly for opening hours and detailed amenity information. These limitations were acknowledged and taken into consideration during the GR calculation.

Weather Data: Current weather conditions were retrieved from the Open-Meteo API, providing temperature, precipitation, and cloud cover data. Weather states were simplified into three categories: sunny, cloudy, or rainy. This condition influenced the distance decay calculations in the relevance assessment.

Physical Environmental Classification: The environmental context utilised the Swiss Federal Statistical Office (BFS) Gemeindetypologie 2020 dataset at a medium precision (G1). This classification distinguishes urban (Type 0), rural/small town (Type 1), and nature-dominated (Type 2) environments (Fig. 6). The data were uploaded as a hosted feature layer on ArcGIS Location Platform, providing environmental context via spatial API calls.

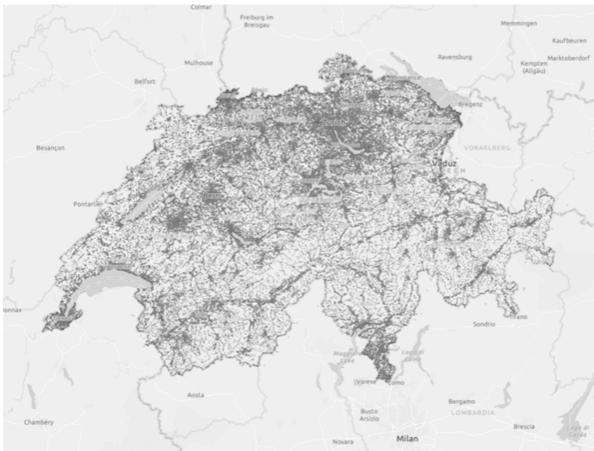

Figure 5: OpenStreetMap point of interest data for Switzerland

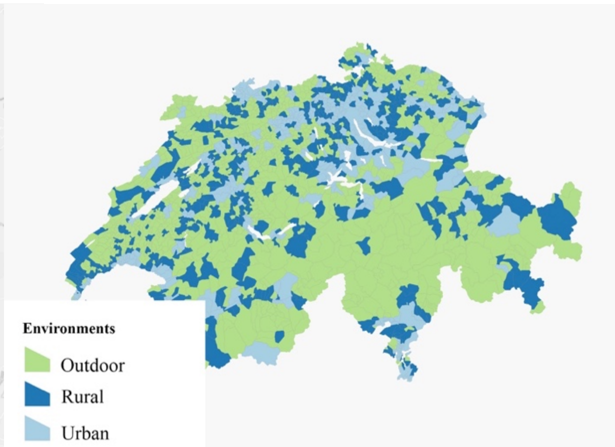

Figure 6: Physical environmenl data of Switzerland green = outdoor; dark blue = rural; light blue = urban

3.3.2 Thematic Modelling

User activities cannot be directly observed and need to be inferred from contextual cues. The thematic preference system introduced a scoring model that combined physical environment context, temporal patterns, and activity-specific probability distributions. These distributions were derived from the previously introduced empirical research predicting user interests across six thematic categories: shopping, food, public transport, culture, outdoor, and explore (a neutral fallback or discovery mode).

The model implemented a weighted noisy-OR framework, a probabilistic approach in which multiple independent factors can each increase the likelihood of an activity being relevant.

Unlike a strict logical OR, where a single cue would be enough to determine the outcome, the noisy-OR allows each cue to contribute with its own probability while also incorporating uncertainty. This approach reflects the reality of multiple contextual cues that can suggest activities independently from each other and all with different degrees of confidence. This framework could weigh the factors individually, while introducing a controlled amount of noise to account for real-world uncertainties.

3.3.2.1 Physical Environmental Base Scores

The algorithm introduces base scores that served as multipliers for activities. These are dependent on the three physical environmental classes from the BFS:

- Urban environment: Emphasised shopping (1.8x), public transport (2.0x), and culture (1.5x), reflecting the concentration of these amenities in cities
- Rural/small town: Balancing all activities with moderate outdoor preference (1.6x)
- Nature environments: Strongly favours the outdoor activity (2.5x) and exploration themes (2.0x).

3.3.2.2 Temporal Activity Patterns

In addition to the physical environment, temporal patterns were shown to also play a central role in shaping activity preferences. Building on the research of Teixeira et al. (2024) and García (2017), these patterns were modelled to capture how activity likelihoods change throughout the day. For instance, food preferences were modelled with a trimodal Gaussian distribution to represent mealtimes. This distribution was chosen to generate fuzzy-edge cases and account for individual behaviour patterns. During the week, peaks were generated in the morning at 08:00, at lunch at 12:00, and in the evening at 19:00 with standard deviations $\sigma = \{1.3, 1.6, 2.0\}$. On the weekend, the food-related activities shift later and were therefore set at 09:00, 12:30, and 20:00 respectively. On the weekends, a wider standard deviation of $\sigma = \{1.3, 1.6, 2.5\}$ was chosen to reflect more flexible dining patterns.

Commuting patterns were modelled using bimodal distributions, with peaks in the morning and evening to reflect public transport demand (Hatcher & Mahmassani, 1992; Li et al., 2004). Urban areas were designed with a stronger peak (scale = 1.2) compared to rural areas (scale = 0.6), reflecting differences in transit dependency.

Shopping patterns were also represented with a Gaussian distribution and mainly varied by day type (Hirsh et al., 1986). On weekdays, a bi-modal distribution was created with a lunch-break peak (12:30, amplitude = 0.8) and a post-work peak (17:00, amplitude = 1.4). On Saturday, a single broad peak was modelled at 13:00, with an amplitude of 1.6 and $\sigma = 2.5$ hours. On Sunday, a reduced morning peak at 11:00 was used with a small amplitude of 0.9 to reflect the limited trading hours and rest-day for a lot of businesses.

Outdoor activities were designed with conditional bonuses that depend on the time of day and factors of the physical environment. Those activities typically peak in the early afternoon and are more frequent on weekends (Haswell et al., 2014; Hayir-Kanat & Breuste, 2020). The activity coincides with daylight availability, which was assessed through a simple Gaussian-centred distribution at 14:00 with $\sigma = 4$ hours.

Cultural interest was modelled with conditional score increments that focused on weekends or evenings (García, 2017). Exploration, by contrast, functioned as a fallback theme that encourages diversity and supports the non-routine behaviour. This theme received a base bonus of 0.8 and was most likely to be selected during early mornings on weekends. Its role reflects theoretical arguments that diversity in recommended content enhances usability and prevents users from being restricted to only a few categories (Chen, Yu, Tang, He, & Zeng, 2017; Liu, Fu, Yao, & Xiong, 2013). The final theme selection happened through maximum likelihood with a Gaussian noise added to introduce a more realistic variability:

```
predicted_theme = max(final_score(theme) + N(0, 0.1))
```

3.3.3 Relevance Modelling

The relevance modelling operationalised the GR framework of De Sabbata (2013) for its optimal balance of theoretical rigour and computational practicality. The CPL provided clear mathematical formalisation for code implementation and allowed the separation of mandatory and desired factors. The following modifications were applied to make the introduced framework fit the mobile context of the prototype application. Most significantly, the model included an assessment strategy for user preferences inspired by RecoMap (Ballatore et al., 2010), where POI clicks were recorded and decayed in significance over time. Favourite POIs were used to create persistent anchor points in personal geography (Couclelis & Tobler, 1987). However, the factor of directionality was excluded since it cannot be reliably computed for stationary or slowly walking users.

3.3.3.1 Geographic Environment Scoring

In contrast to the physical environment that was used in the thematic modelling process, two additional geographic factors extended the framework by quantifying spatial relationships that were not yet captured by the traditional POI attributes.

As introduced in the theoretical part, the **cluster score** quantifies how densely a POI was surrounded by similar establishments, combining cluster membership and proximity to nearest neighbours. This score identified natural commercial districts where the concentration of features enhanced individual relevance. With the help of density-based spatial clustering, the score was assessed with the following logic:

```
Cluster_score = √((size_component² × distance_component²) / 2)
```

The cluster size (size component) reflected the relative cluster size, and the distance captured the proximity of the entity to the cluster centre. The first reflected that a larger cluster should be more relevant, and the latter stood for the density of the cluster. As De Sabbata (2013) proposed, essential services (e.g., hospitals or schools) automatically received maximum scores, as they do not follow commercial clustering patterns and therefore fulfil the clustering property.

The **co-location score** identified beneficial proximity patterns through the spatial association rule mining. The algorithm tried to detect spatial association rules and results in rules such as 'tourist information → viewpoint'. This was carried out by analysing POI co-occurrences within a 200-meter radius. High-quality rules with a minimum support of 30% and a lift larger than 1.5 were applied to score POIs based on proximity to complementary amenities:

```
Colocation_score = mean(top_rule_scores)
```

Both scores were pre-computed using a custom Python script that connected to ArcGIS Location Platform and processed the large number of POIs using KD-trees, which are widely shown to support O(log n) query time for exact match and nearest-neighbour searches. The code directly updated the hosted feature layer and added two new attribute columns with the computed scores. This pre-processing shifted the intensive spatial analysis from runtime to initialisation.

3.3.3.2 Relevance Score Computation

The CPL framework used generalised conjunction or disjunction (GCD) and conjunctive partial absorption (CPA) operators. The GCD operator with $\alpha = 0.75$ implemented soft conjunction, requiring reasonably high values across the inputs without harsh cutoffs to also result in a high output. The CPA operator then combined mandatory requirements with desired enhancements.

Semantic/Topicality Component (Mandatory)

The semantic component performed a binary matching between the POI amenity type and the predicted thematic interest from the TM:

```
P_semantic = {1.0 if fclass ∈ Theme_fclasses, 0.0 otherwise}
```

The binary approach ensured a lower relevance score for POIs that were not thematically aligned with the modelled user interest or activity.

Spatio-Temporal Proximity (Mandatory)

Spatial distance and temporal availability were used to model the availability and physical proximity of a place. Distance decay aimed to model human spatial behaviour where relevance dropped only gradually at nearby places and decreased more rapidly further away:

```
Base_distance = {
    1.0 km if speed < 7 km/h (walking)
    3.0 km if 7 ≤ speed < 15 km/h (cycling)
    5.0 km if speed ≥ 15 km/h (driving)
}

Weather_factor = {
    1.0 for sunny (weather = 1)
    0.7 for cloudy (weather = 2)
    0.3 for rainy (weather = 3)
}

σ = Base_distance × Weather_factor
P_distance = exp(-0.5 × (distance / σ)²)
```

Weather adjustments accounted for the situation that users were less likely to travel longer distances in bad weather, favouring places more closely located to their current location.

Temporal availability was parsed from the attribute string, checking against current time and day. The non-zero score for closed POIs allowed for planning ahead while strongly preferring currently accessible places:

```
P_temporal = {0.6 if POI is open, 0.1 if POI is closed}
```

Behavioural Factors (Mandatory)

User interaction history created personalised relevance through two mechanisms:

```
P_favourite = {1.0 if marked as favourite, 0.3 otherwise}
Click_rate = click_count / (days_since_last_click + 1)
P_clicks = 1 - exp(-click_rate / 5.0)
P_behavioural = max(P_favourite, P_clicks)
```

Favourites acted as persistent anchor points, while click-based preferences decayed over time. The exponential formulation ensured that recent interactions with the feature carried more weight.

Geographic Environment (Desired)

The pre-computed cluster and co-location scores were simply retrieved and combined:

```
P_geographic = GCD_0.75(cluster_score, colocation_score)
```

The final relevance computation followed the following designed hierarchy and began with the spatiotemporal fusion:

```
P_spatiotemporal = GCD_0.75(P_distance, P_temporal)
```

Initial mandatory combination:

```
P_mandatory = GCD_0.75(P_semantic, P_spatiotemporal)
```

Behavioral integration with stronger conjunction (α=0.85):

```
P_mandatory_final = GCD_0.85(P_mandatory, P_behavioral)
```

Final relevance through CPA:

```
Relevance_score = CPA_0.75,0.75(P_mandatory_final, P_geographic)
```

The stronger conjunction with $\alpha = 0.85$ for behavioural integration was chosen to ensure that user preferences significantly influenced relevance. A favourite restaurant, for example, maintained a high relevance even at a greater distance. The resulting scores of the CPA were in [0,1] and reflected a nuanced relevance relationship.

3.3.4 Model Training

Using the rule-based relevance and thematic scoring functions, synthetic training data was generated to make the relevance assessment based on ML-based predictions. The data generation process created realistic user-POI interaction scenarios through controlled randomisation. As an output, the thematic dataset comprised 1’000 samples and 8’000 samples for the relevance assessment. Each sample contained a full context vector and a computed ground truth score. Afterwards, these datasets were enhanced with noise augmentation to improve generalisation during model training. Continuous features were perturbed with Gaussian noise. For categorical features, random flips were applied to simulate occasional misclassifications. Temporal variables were shifted within a ±3-hour window to reflect variability in activity timing. Finally, binary features were altered through bit flips to capture uncertainty in yes/no inputs. The aim of these augmentations was to simulate real-world noise in sensor readings and user behaviour.

Both models were trained with XGBoost (eXtreme Gradient Boosting) due to the algorithm's ability to handle mixed data types, missing values, and complex non-linear relationships (Chen & Guestrin, 2016). In particular, it supports heterogeneous features (continuous, categorical, binary) and is robust to incomplete data. The latter was crucial when working with OSM data due to the inconsistent nature given in the publicly sourced attributes.

Despite its proven effectiveness in spatial applications (Yu & Xie, 2024; Zhu, Shu, & Zou, 2022), XGBoost also presented several challenges. The computational overhead was managed by constraining the model to a maximum of 300 trees with a depth of 4. To reduce the risk of overfitting, the parameters were regularised aggressively, while the issue of limited interpretability was addressed through analysing feature importance outputs.

Prior to training, the relevance model further incorporated monotonic constraints to enforce logical relationships:

```
monotone_constraints = {
    distance: -1,       # ↑ distance → ↓ relevance (always)
    isOpen: 1,          # open POI → ↑ relevance
    favorite: 1,        # favourite → ↑ relevance
    clickCount: 1,      # ↑ interactions → ↑ relevance
    clusterScore: 1,    # ↑ clustering → ↑ relevance
    colocationScore: 1  # ↑ co-location → ↑ relevance
}
```

The TM model, however, applied no constraints, as the activity–context relationships were complex and highly context dependent.

The training strategy was divided into two stages. The first stage included a hyperparameter optimisation with a five-fold cross-validation in 80% of the augmented data. The code then searched through over 700 and 400 combinations of parameters for the RM and TM, respectively, to find the optimal parameters for model training. Afterwards, the second stage trained the final model on 100% of the dataset using the previously detected optimal hyperparameters. Finally, the two ML models were exported and saved so that MyPlaces could use them during operation for GR assessments. Figure 7 provides an overview of the overall GR modelling strategy, illustrating how the different attributes and variables interact in relation to GR, GIN, and GIOs.

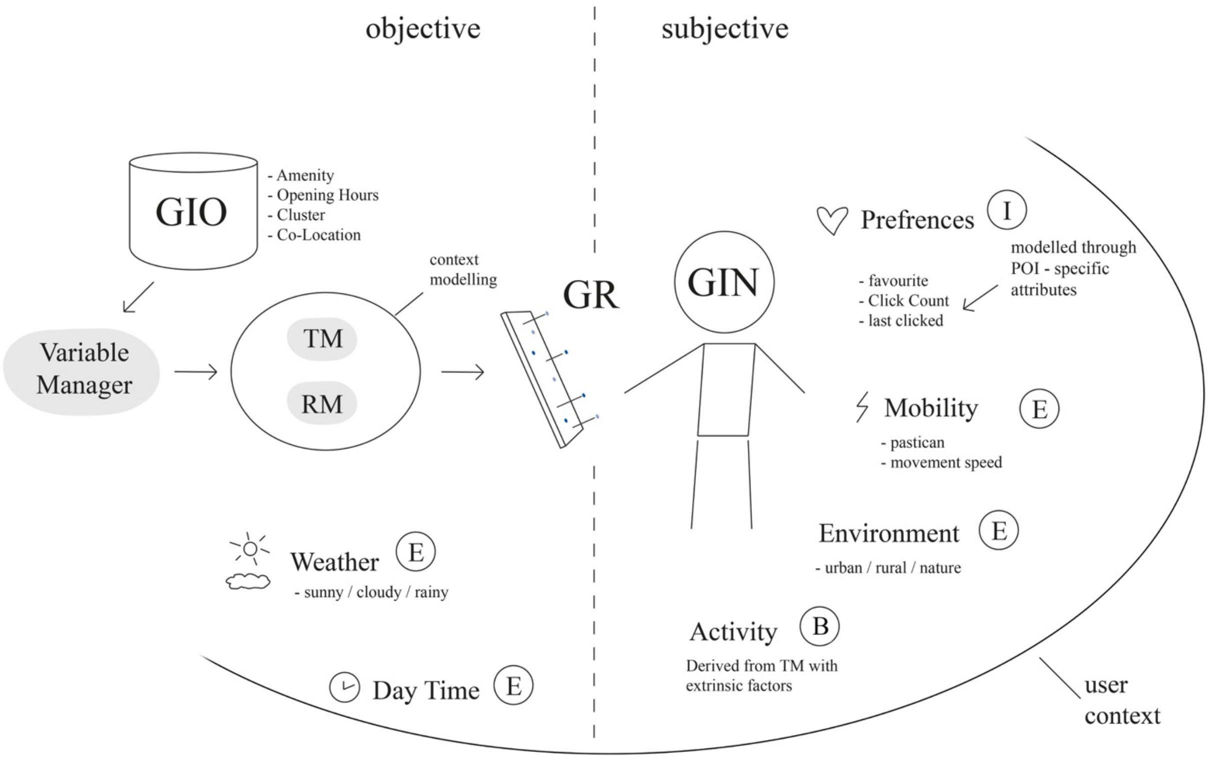

Figure 7: Context modelling strategy in relation to the different influential context attributes

3.4 Visual Design Strategy

The visualisation strategy of the mobile map application was designed to communicate relevance intuitively and effectively, especially under real-world mobile usage constraints. Building on established cartographic design principles, the system employs a combination of visual variables to encode the relevance of POIs. The visual interface of the application was designed to communicate POI relevance visually and to provide a minimal, intuitive, and interactive user experience aligned with HCI guidelines. The goal was to reduce cognitive friction while maximising spatial awareness and relevance comprehension.

The map forms the visual and functional core of the application. It is a spatial representation of the GIOs and a dynamic layer for communicating GR to the user. Drawing from the established HCI and cartographic design principles that were introduced in Chapter 2, the map was designed to support quick interpretation, intuitive interaction, and adaptive content filtering.

3.4.1 Relevance Based Points of Interests

The central mechanism for highlighting the relevant POIs was grounded in visual variable encoding. Informed by the presented prior research on effectiveness in mobile contexts, the application employed the three most effective primary visual variables to communicate the GR: size, saturation, and transparency. Size encodes the quantitative relevance score of a POI since larger symbols are perceived as more important (Cláudio et al., 2011; Förster, 2023). Saturation and transparency refine this hierarchy: vivid, fully opaque icons emphasise relevance, while desaturated and semi-transparent ones recede visually into the background without being removed. Less relevant POIs only appear at closer zoom levels, ensuring that the map remains uncluttered yet complete. This approach balances clarity and control, allowing users to focus on the most relevant features while still accessing all POIs when needed.
The colour hue was used to denote the grouped entity type of the feature (e.g., food, shopping, culture). The colour palette was aligned with familiar schemes from established mapping applications to leverage existing mental models and foster intuitive comprehension (John, 2014). Analysing the OpenStreetMap GitHub Library,[3] healthcare appears red, transportation blue, landmarks grey, and food orange. Overall, there is not a lot of hue variation, and the chosen markers appear rather monochrome. Apple uses yellow for shops, pink for cultural amenities, and light blue for sport-related places. However, there is no uniform standard defined. An overlap between the standards of OSM and Apple Maps was chosen to guarantee a harmonised colour choice for the iOS user.

To enhance readability across varying basemap conditions, POI symbols were designed with white outlines, which is a choice supported by tourist map symbol studies showing that framed pictographic symbols improve legibility (Konstantinou, Skopeliti, & Nakos, 2023). The symbology employs Apple's SF Symbols library to ensure visual consistency with the iOS ecosystem. Their integration ensures that icons are visually cohesive, homogeneous, and technically optimised for various device resolutions and accessibility features.[4] These pictographic symbols enhance recognition by resembling real-world objects, reducing search times and improving detectability (Forrest & Castner, 1985).

[3] https://github.com/gravitystorm/openstreetmap-carto/tree/master/style, accessed 19.09.2025.
[4] https://developer.apple.com/design/human-interface-guidelines/sf-symbols, accessed at 15.09.2025.

Map overcrowding by assessing many features as relevant to the user was solved with an aggregation mechanism, which is explained in greater detail in Section 3.5.3. A slightly increased density of POIs compared to traditional digital maps was favoured in the visualisation, since it has been proven to increase user comfort and task success (Bartling et al., 2021).

3.4.2 Basemap Design

The basemap plays a crucial supporting role by providing spatial context for the POIs while minimising visual distractions. The topographic basemap style from Esri's ArcGIS Location Platform was selected as the foundation, as similar styles have demonstrated superior performance, task success, and confidence compared to other styles (Bartling et al., 2021). Three basemap variants were designed to support different usage contexts:

- Day mode: Based on the topographic style with adjusted land transparency and hue to ensure readability and seamless visual integration.
- Night mode: Derived from the street style with streets and tunnels transformed to grey hues, supporting the visual hierarchy of the POIs.
- Food Theme: Uses the same colour hues as food-related POIs and applies warmer tones overall to create a cosy, inviting atmosphere suited for dining contexts.
- Transportation Theme: Highlights public transport routes and lightly emphasises street networks to aid orientation, while avoiding visual dominance over POIs.
- Outdoor mode: An enhanced topographic day mode variant that featured additional contour lines for elevation context.

All styles were adjusted in the Vector Tiles Styles Manager from Esri. The customisation strategy involved removing transportation POIs from the basemap layers (as these were managed through the dedicated POI feature layer) and adjusting colour values to maintain consistency across all the map variants. Due to the already constrained screen space availability on mobile devices, classic map elements, such as a legend, scale, or north arrow, were not included (Reichenbacher, 2004). Example extents of the five designed basemaps can be found in Appendix 1.

3.4.3 User Interface Design

The UI overlay architecture followed iOS interface guidelines while implementing Gestalt psychology principles. The design employs a functional grouping strategy, organising interface elements into two distinct areas that can be seen in the left part of Figure 8. The top area includes the map configuration tools with the search bar, a theme picker, and a dark mode toggle. The lower part of the screen features a personalised 'home' area with controls for the user profile and the favourite POIs. This arrangement leverages the principle of proximity to create perceived order while maximising the central map viewing area.
The visual design language utilises iOS's glass morphism approach and uses blur effects to create depth through translucency, allowing the map to stay partially visible while maintaining sufficient contrast for legibility.

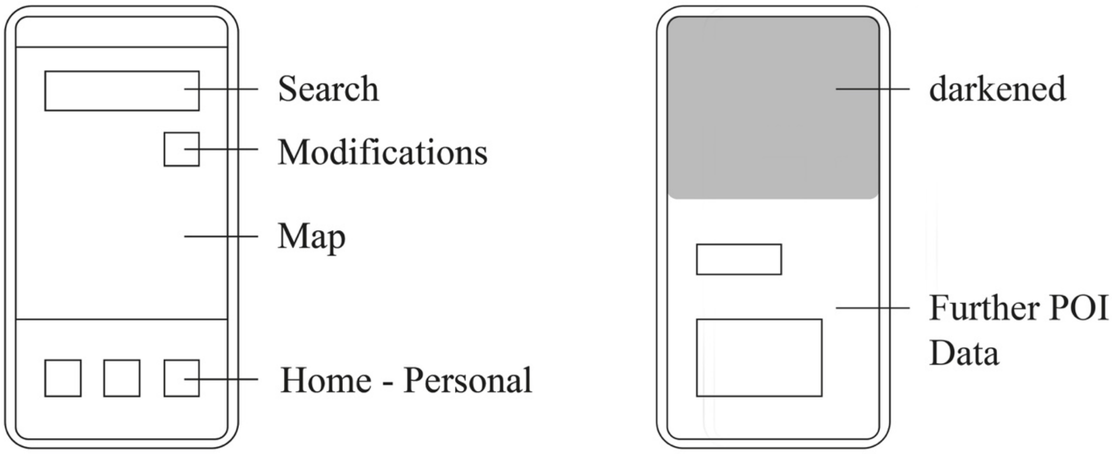

Figure 8: Overview of the user interface design strategy and layout

The colour encoding followed platform conventions, with blue (systemBlue) consistently indicating interactive elements[5] – a signifier that affords manipulation. This semantic colour mapping extends to state changes: a green checkmark for active users, yellow stars for favourites, and subtle grey stars for disabled stars. Typography employed San Francisco (SF) fonts with deliberate hierarchical scaling. Primary labels use .title2 and transition into regular .caption for secondary information. The differentiation helps create a clear visual hierarchy. Each UI element features drop shadows that established perceptual layering, making overlays appear to 'float' above the map surface while maintaining sufficient separation from the basemaps.

[5] https://developer.apple.com/design/human-interface-guidelines/color, accessed 24.09.2025.

The interaction design implemented visual feedback through animations and state transitions. Touch interactions trigger minimal-scale transformations and opacity changes that confirm the user input. Spring animations were used to provide natural physics-based transitions that aligned with users' mental models of object manipulation. This animation was applied to all expandable items and was marked with respective chevron icons or drag indicators to signify the affordance to expand.

When a new view is entered or expanded, the hierarchy principle brings it to the foreground while keeping the original view visible in the background to support orientation. This idea of modality was implemented by partially darkening or overlaying the old view with the new one, as illustrated in the right panel of Figure 8.

3.5 Architecture and Component Design

This section outlines the software architecture and core components of the AMMA MyPlaces. As proposed by the theoretical background, the system was designed following a modular layered architecture that separated concerns, such as context handling, relevance scoring, and data persistence. The following subsections describe the modules, their interactions with each other, and the design patterns implemented in the system.

3.5.1 Architecture Overview

The application follows a three-tier architecture with distinct layers for presentation, logic, and data management. At its core, the system includes multiple specialised components that work together and deliver adaptive behaviour to the user. An overview is visualised in Figure 9, and the detailed class diagram can be found in Appendix 2.

To handle the geographic data, the ArcGIS Maps Software Development Kit (SDK) was used in combination with ArcGIS Location Platform data to customise the application. It features possibilities for a large number of spatial operations and data hosting. ArcGIS Maps SDK for Swift then enabled the mapping and GIS capabilities for iOS, such as adding feature layers, geolocation, feature querying, or even routing.

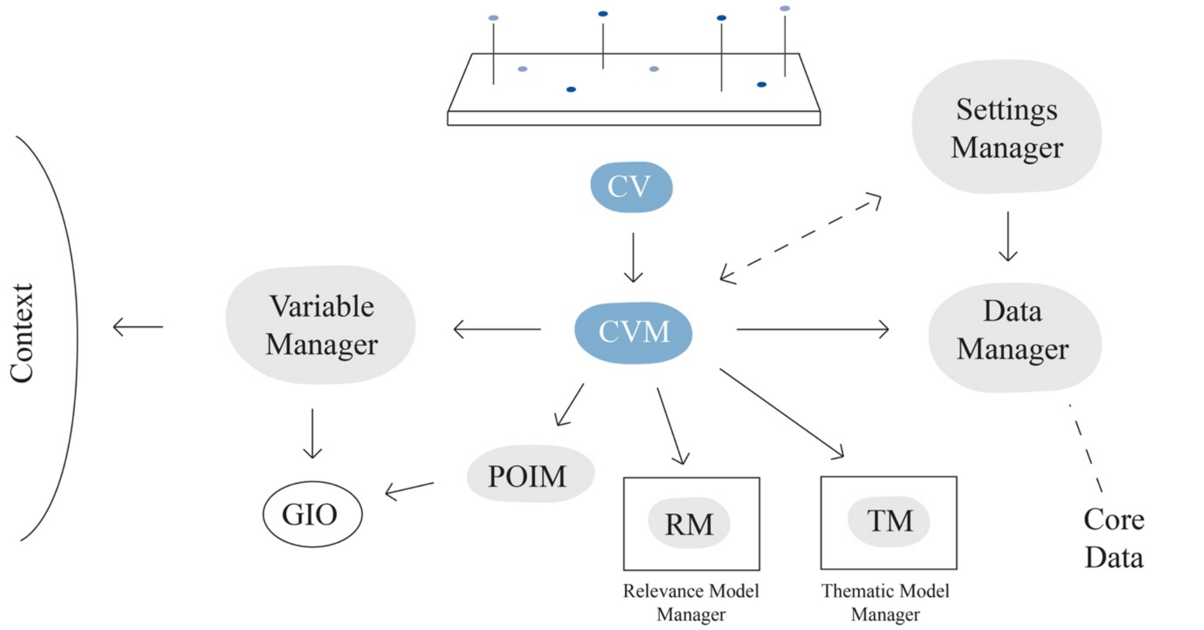

Figure 9: Overview of the most relevant components of the system architecture

3.5.1.1 Core Components

The content view model (CVM) is designed to be the main orchestrator managing the adaptation logic and coordinating between the other system components. It handles user interactions with the UI, POIs, and the map, processes location changes, and triggers relevance computations. The model hosts the loaded POI through the lifecycle and filters and aggregates them until the final presentation. It also holds the different basemap styles and the additional POI layer copy for the visualisation of the irrelevant features.

The variable manager acts as the context computation engine, which is responsible for collecting and processing all the contextually necessary inputs for the RM and the TM. It communicates and connects with external services, such as the weather API from Open Meteo,[6] the hosted physical environment data via API call from the BFS or the device sensors (location, time). It maintains a cache to reduce the API calls and converts the raw POI attribute values in combination with the additionally gathered data to converted integer inputs for the TM and RM. The relevance model manager is the encapsulation of the RM engine, wrapping the core ML model for relevance prediction. It provides a clean interface for the CVM to hand over the 11 attributes for the GR assessment and handles the model initialisation, input pre-processing, and output interpretation.

[6] https://open-meteo.com, accessed at 21.09.2025.

The same is true for the thematic model manager, which wraps the TM and also provides a simple access point for the CVM to predict the theme. The output is one of the six pre-defined possible thematic interests of the user.

The POI model fetches and handles the POI data from ArcGIS Location Platform. It implements a bounding box fetching to efficiently retrieve and load the relevant POIs around the user location. The fetched features are then passed to the CVM for further processing. The buffer size was chosen so that the number of fetched POIs approach but never exceeded the fetching limit of 2,000 features. Therefore, arbitrary sub-sampling was prevented so that no important POI could be excluded.

A core data persistence layer is provided by the data manager. It handles the storing and retrieving of user profiles, relevance scores, and interaction history and manages the relationships between the users, POIs, and the assigned scores. It provides methods for the management of user favourites and user switching.

Finally, the user preferences and application state are handled with the settings manager, which ensures consistent access across the application. It manages theme selection (manual and predicted), basemap switching, and user profile states for the data manager.

3.5.1.2 Data Model

The application implements an on-device persistence layer using Apple’s core data framework. This ensures user privacy while maintaining personalised experiences across multiple profiles. The architecture consists of the following components:

Persistence controller, which manages the core data Stack; the data manager, which provides a repository abstraction of it; and the core data schema, which defines the relational model. The core data schema implements a many-to-many relationship between users and POIs, mediated through the relevance score entity, as visualised in Figure 10.

The user profile entity stores user identification, personal information, and thematic state. Each profile has a one-to-many relationship with relevance score entities, enabling multiuser support with a complete data separation.

The POI entity provides references to geographic features with its own identifier and a specific fid value that allows synchronisation with the external ArcGIS feature service.

The relevance score is the associative entity matching the POI and their user-specific attributes to the correct user. It stores foreign key values in combination with user preference data (isFavorite, ClickCount, lastClickedDate), the relevance scores as well as the context attributes used to predict the relevance score.

This data architecture guarantees complete data sovereignty, and all personal information and the interaction history remain stored locally on the device in an SQLite database. The application only fetches generic POI data from external services and computes the relevance score on the device with the ML models and never transfers user-specific data externally.

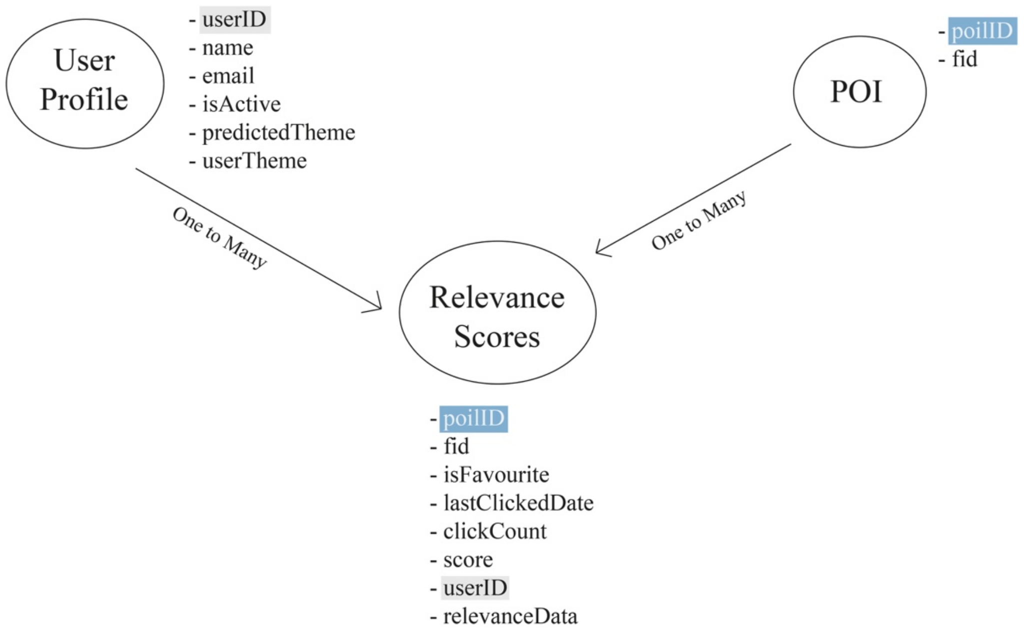

Figure 10: Core data entities with their relations and variables (identifiers (IDs) for joins are highlighted in the same colour)

3.5.2 Design Patterns

The application architecture makes use of several design patterns to guarantee extensibility and maintainability. The most important logical implementations are:

- *Singleton Pattern:* The settings manager and data manager are initialised at application launch and, through the Singleton pattern, guaranteed global access to data retrieval and storage as well as settings modification or retrieval. It guarantees a global point of access to that one single instance of the two objects.
- *Observer Pattern:* This pattern enables reactive UI updates, such as changes in relevance scores, theme switches, or favourites. It is enabled through the @Published property and allows for informing the dependent objects about an update to let them react accordingly.

- *Repository Pattern:* The data manager implements this pattern to abstract the core data operations behind a clean interface for the other objects to access.
- *Adapter Pattern:* The variable manager acts as an adapter, converting heterogenous context sources (weather data, location data, attribute data) into unified integer outputs for the scoring system.
- *Facade Pattern:* The content view model provides a simplified façade for the complex subsystem of the POI loading, scoring, and aggregation, only exposing high-level methods that are crucial for the application.
- *Cache Pattern:* Cache pattern is implemented in the variable manager, such that strong expensive computations (weather and physical data retrieval) are stored in the cache with 15-second validity windows to minimalise the API calls while computing the GR and still maintaining data freshness.

3.5.3 Application's Data Flow

The system's data flow follows a pipeline architecture with distinct phases that are visible in a simplified version of the State Diagram in Figure 11. A clear overview of the State Diagram is in Appendix 3. At the application startup, the system loads the user profile from core data and chooses the one that was active in the latest session. If there is no user created yet (first app launch), a special onboarding view shows up to create the first initial user profile. The user is redirected to the content view (CV). As soon as the CV shows up, the POI model (POIM) gets initialised, and location monitoring begins. The POIM queries the features from ArcGIS with spatial and semantic constraints to exclude amenities that are not assumed relevant at any point (automatic teller machines, vending machines, etc.). Features having no name attribute are also filtered out. The POI retrieval is limited with a bounding box around the user's location, as discussed in Section 3.5.1.1.

After the initialisation, the application starts to collect contextual data to predict the thematic interest and the relevance for the loaded POIs. The CVM starts the modelling process and lets the variable manager orchestrate context gathering, caching expensive operations and deriving attributes for correct inputs in the two prediction models. The computationally expensive operations (weather API calls, physical environment retrieval) that are used a lot of times are retrieved only once per relevance score update. With a 15-second time validity, the data freshness is preserved for an eventual context change soon after while still minimising network load. The variable manager computes derived metrics, such as distance, using the efficient spatial calculation from the geometry engine library with the correct coordinate system handling. For coordinates that are already in WebMercator projection, the simple Euclidean distances suffice. Geographic coordinates (WGS84, LV95) leverage iOS's native CLLocation methods and get cached to improve efficiency. The variable manager derives all the temporal variables from the device's date and time information.

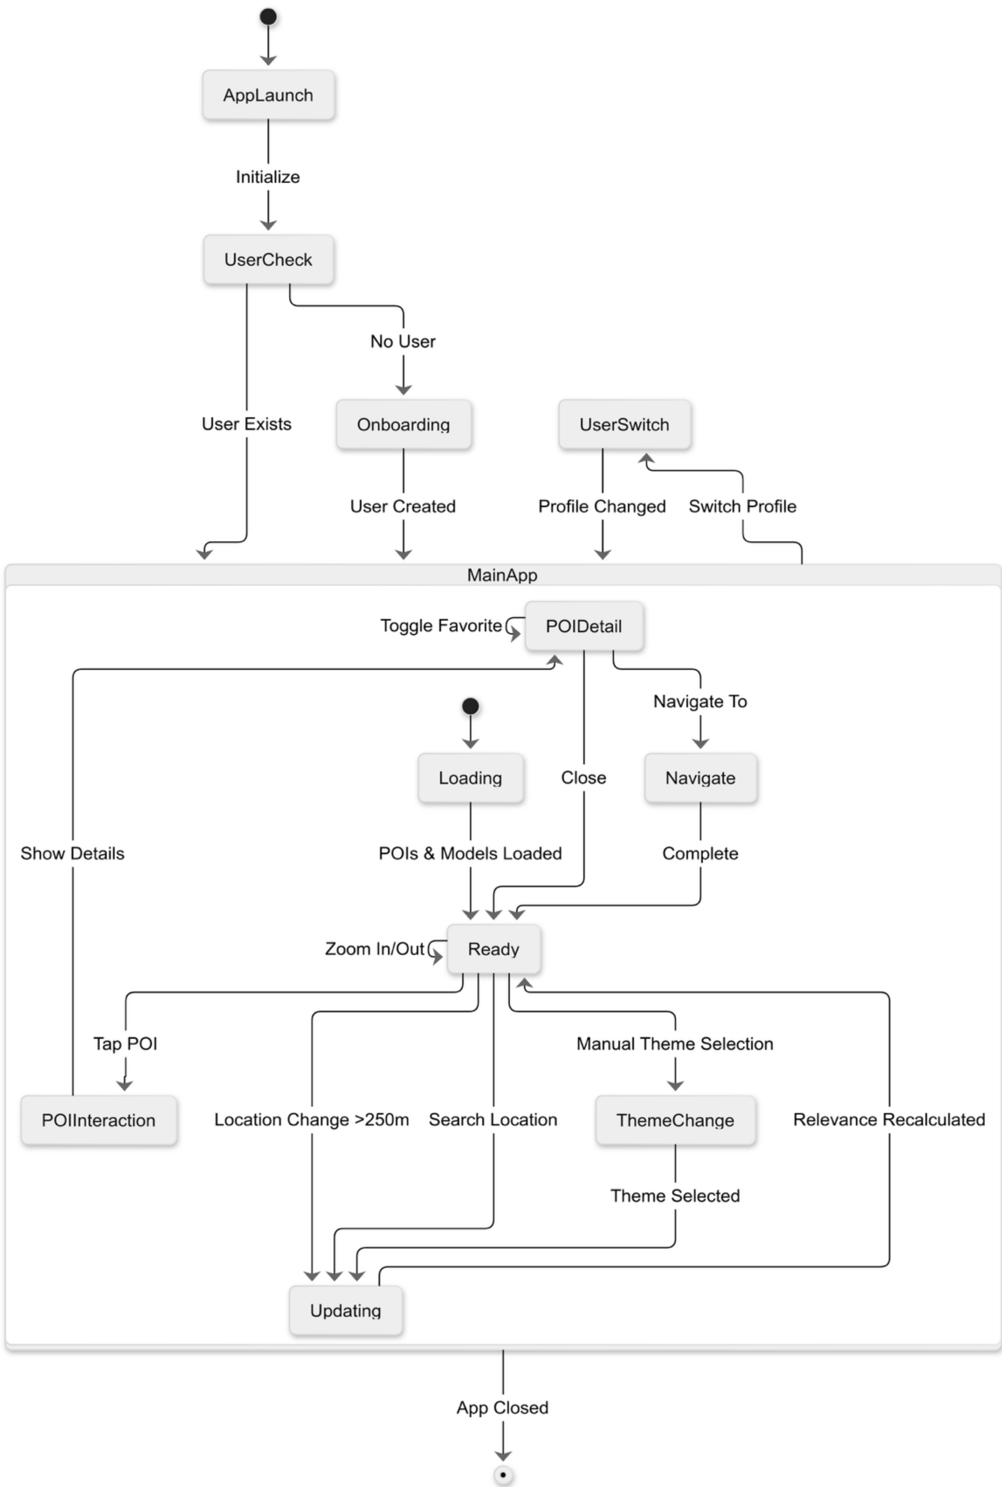

Figure 11: Simplified state diagram of MyPlaces

The TM continues with the quickly gathered information to set the initial thematic choice. This is then aligned with the correct basemap and published to the CV together with the irrelevant POI feature layer. Concurrently, as soon as the active user is loaded, their favourite POIs are loaded to the favourites panel in the UI through background processing that does not block the main thread.

For each loaded POI in the POIM, the system computes the GR score using the ML model. The scoring pipeline considers the semantic alignment between theme and POI amenity, the spatio-temporal proximity, the interaction history, and its geographic environment pattern. The POIs are then filtered and added to the CV. In this process, the TM and RM work in tandem to create an adaptive experience. The TM provides semantic grouping and categorises the POIs into contextually relevant themes based on the temporal and physical environment information. The RM ranks and personalises within and across themes, considering the individual POI characteristics and user preferences. Together, they transform the static map into a dynamic interface.

As soon as the user interacts with the map and changes the scale more than 25% smaller or larger in comparison to the last aggregation scale, the CVM newly filters all the POIs to prevent cluttering. The model evaluates the most relevant POI within a pre-defined cluster distance around features. This distance varies according to scale and reaches from 3 m up to several hundred meters. The newly computed layer is then again added to the map.

This operation is handled with a debounce to ensure smoother pinch-zoom interactions without computational overhead. The system maintains two separate arrays for filtered and displayed POIs, as mentioned in Section 3.5.1.1, so that they can be re-aggregated without a new analysis of relevance.

```
function adaptiveAggregate(pois, scale) {
    threshold = getThreshold(scale)
    clusters = []
    for poi in pois {
        nearby = findWithinThreshold(poi, threshold)
        best = selectByRelevance(nearby)
        clusters.add(best)
    }
    return clusters
}
```

Location changes are similarly optimised through throttling: a 250-meter movement threshold filters minor position changes, while a 2-second minimum interval prevents cascade updates during rapid travel. As soon as this limit is exceeded, a new context modelling process with the TM and RM is initialised.

The context-modelling process is also updated if either the user overrides the thematic prediction or if they search for specific places. In the first case, the user can manually override the thematic choice via the UI and cascade the new theme to the RM for further processing. Search operations establish a new virtual location, which additionally involves a new loading of the relevant POI features around the new geographic position.

The user can switch between day and night mode via the CV, which is then linked to the settings manager. There, the current basemap state is stored. Upon application launch, the CVM initialises the map style based on system time: if the time falls between 08:00 and 19:00, a light basemap is loaded; otherwise, a dark basemap is applied. When the user changes the state, it notifies both the CVM and settings manager, resulting in an immediate basemap switch and view update.

User switching invalidates all the cached data and favourites, loading the new user's preferences and interaction history before triggering complete relevance recalculation. This ensures personalised experiences remain isolated between user profiles.

3.6 Evaluation Approach

This section outlines the evaluation methodology that was used to assess the technical feasibility and performance of the AMMA. Given the prototype nature of this implementation and the focus on demonstrating technical viability, the evaluation strategy emphasises quantitative performance metrics, model accuracy assessment, and computational efficiency rather than user studies.

The evaluation approach followed a multi-faceted strategy to validate both the ML components and the overall system performance itself. Rather than conducting user studies to measure usability of the application, which would require extensive participant recruitment and controlled experiment settings, this thesis focused on establishing technical feasibility through systematic performance measurement and model validation. This approach enabled rapid iteration and provided quantifiable metrics that could inform future research.

The evaluation framework comprised two primary components: ML model validation using the generated synthetic relevance data and testing under noise conditions, and runtime performance measurement in realistic deployment scenarios.

3.6.1 Model Validation Methodology

The synthetic data produced through the rule-based CPL in Section 3.3 served as the ground truth for training and evaluation. Since the resulting score saved in the dataset was already created with the rule-based functions, a simple analysis would lead to functions exceeding the models by far. Therefore, the performance comparison between the ML models and the rule-based functions was analysed under varying noise conditions (0–30%). The same noise rules as in section 3.3.4.2 applied.

For the regression tasks (relevance scoring), root mean square error (RMSE) was used as the primary metric for the continuous score predictions. Mean absolute error (MAE) was used for interpretable error magnitude and the R^2 score was used to measure explained variance. The classification tasks (thematic prediction) were evaluated through overall accuracy, per-class precision/recall (F1 score) across all six themes, and confusion matrices to identify systematic misclassifications.

3.6.2 Runtime Performance Evaluation

A performance collector captured real-time metrics across the processing pipeline, measuring the POI loading time (feature service query latency); relevance scoring time (ML inference for all POIs); thematic prediction time (context classification overhead); aggregation time (spatial clustering computation); and total pipeline time (end-to-end latency from data request to output rendering). Performance evaluation was conducted in three representative environments at scales of 1:500, 1:1’500, and 1:5’000. The urban case, represented by Zurich, featured a high density of POIs, while the rural setting of Hombrechtikon (Zürcher Oberland) showed a medium density. In contrast, the outdoor environment of Oeschinensee contained only a sparse distribution of POIs. These variations allowed target performance thresholds to be analysed in relation to the HCI guidelines presented in Section 2.2.1.1.

Some internal variables were made publicly accessible to the performance collector, which temporarily violated the OOP principles. This trade-off was necessary for benchmarking and precise performance measurements.

4 Results

This chapter presents the outcomes of implementing the AMMA MyPlaces, demonstrating how the theoretical frameworks and the methodological approaches translate into a functional system for GR assessment. Building upon the visual and technical design strategies outlined in Chapter 3, the results consider both the technical feasibility and practical effectiveness of real-time context-aware POI recommendations on mobile devices.
This chapter considers the three key areas: the application outcome, showing the realised interface and interaction patterns, the ML performance, and the system benchmarks.

4.1 Application Outcome

The implementation of MyPlaces demonstrated the theoretical considerations in practice. The following section presents the realised interface components, interaction patterns, and GR communication as they appeared in the final product.

4.1.1 Final Application Overview

The application comprised five primary views that supported the defined use cases and requirements from onboarding up to daily usage or scientific research. The most important four of them are visualised in Figure 12.

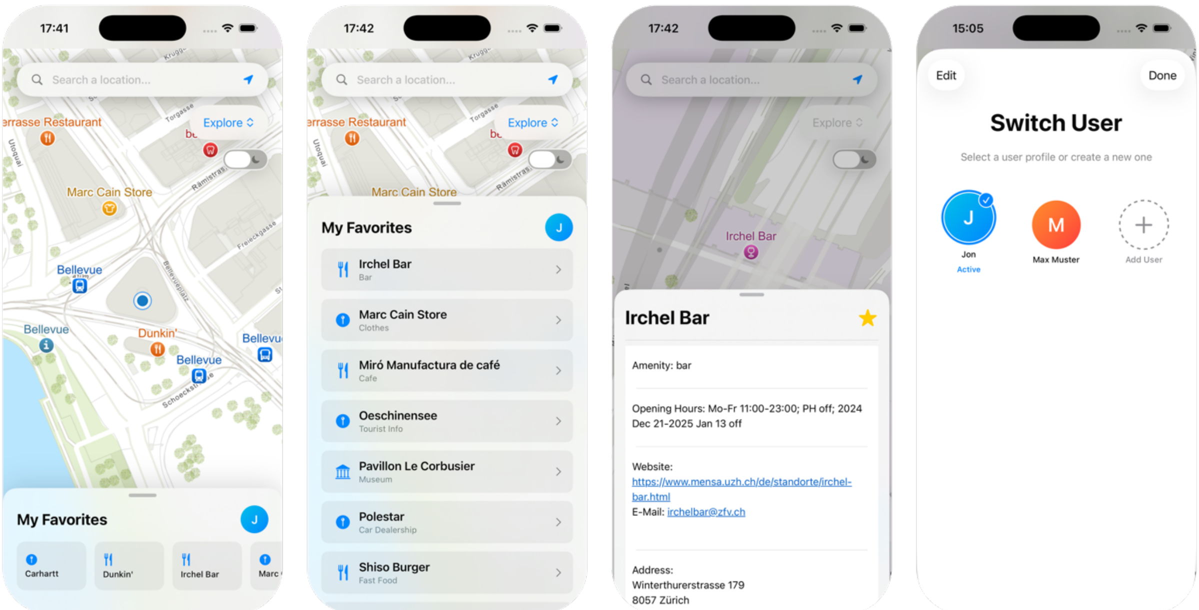

Figure 12: Screenshots of the four main views from MyPlaces
From left to right: Map view, favourites view, popup view, and user profile view

The **map view** serves as the primary interface for spatial exploration and POI interactions, featuring a full-screen map display with overlaid UI controls. It includes a prominently placed search bar for location lookup, the ability to reposition the map extent back to the current user location, and a clear indication of the user's real-time position, which is marked with a circular blue marker and a surrounding halo.
The map presents nearby relevant places and integrates a set of supplementary features to enhance manual map modifications. For instance, a thematic picker allows the user to see and modify the modelled thematic interest of the modelled activity. It aligns with the pill-shaped form of the search bar in a more compact design and is aligned on the right side to create a subtle sense of connectedness with the other interface elements at the top. Located right below this picker, there is the toggle for the day and night mode modifications.
The favourites panel anchored at the bottom of the screen provides access to pinned locations and contains a profile icon for managing user settings. A breakdown of the components and respective layers is provided in Figure 12.

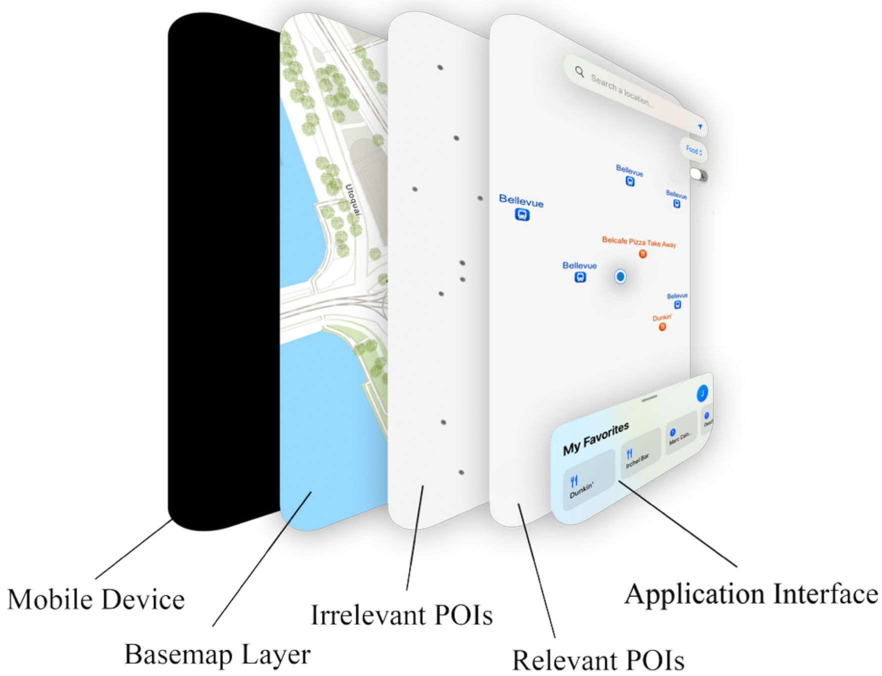

Figure 13: Layer structure overview of the application

From the map view, the user can access three other views. By swiping up the bottom panel, the favourites view is expanded, displaying a list of bookmarked locations. Selecting a location, either from this panel or directly on the map, opens the popup view, where the detailed POI attributes are displayed. Tapping the profile icon in the top-right corner of the favourites panel leads to the user profile view, enabling users to switch between different profiles.

The **favourites view** is anchored at the bottom part of the UI and establishes a reference area with bookmarked POIs and the user profile settings (Fig. 14). The layout makes use of an adaptive layout strategy, presenting the saved POIs through a horizontal scroll view in its collapsed state and transitioning to a vertical list format when expanded. Each favourite element combines a custom icon derived from amenity classifications with the location's name, creating a discrete visual unit that maintain informational coherence. Besides the panel's title, the header section also includes the user's profile avatar, which is generated through a deterministic colour gradient derived from their email hash, ensuring a consistent and recognisable visual identity across sessions. The interface uses visual closure by partially showing the encapsulated POI representation at the right screen border, which implicitly indicates scrollability and the availability of further elements. This design principle is equally applied to the expanded list view at the bottom of the screen. In the case of an empty state, a placeholder message with an instructional subtext is displayed to inform the user that no favourites have been added yet (Fig. 14).

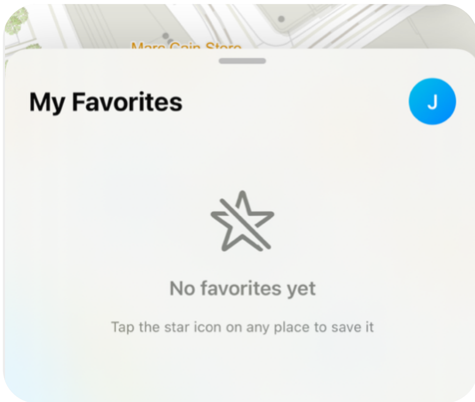

Figure 14: Initial view of the favourite panel with no features yet added

The **popup view** implements a modal presentation pattern that preserves spatial context in the background while focusing user attention on the POI details. It is activated through a tap gesture on the POI marker and presents attributes, such as the name, address, and opening hours of the location. A prominently placed star icon enables the user to add the POI to their favourites and provides immediate feedback when activated. The underlying data attributes are updated to ensure that this favoured place ranks higher in the next GR assessment. Interactive attribute values, such as the website or email, can be clicked and are opened in the device's native browser or email application and make the interaction flow transfer seamlessly outside of the application.

The **profile view** includes user management with a responsive grid layout that scales to include and display all registered users. The active user is visually separated through a blue checkmark, while an 'Add User' card is located at the last position of the grid to provide a possibility to create new profiles. Clicking that, a new modal sheet opens on top and allows entry of the new user details. An edit mode is available for profile deletion with red indicators providing clear feedback during the process. The system enforces a visual constraint to prevent the removal of the active user profile by greying out the profile card and adding a lock symbol (Fig. 15).

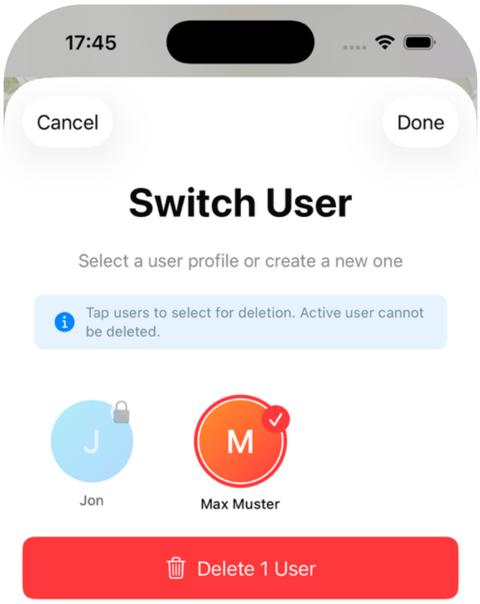

Figure 15: UI of the edit mode while user profiles are deleted

The **onboarding view** adopts a minimalist approach during the initial application startup. It presents only the two essential input fields (name and email) within a vertical layout. On top of that, a bold welcome message is shown, and at the bottom, a blue clickable 'Get Started' button is located that saves the user profile and opens the Map View.

4.1.2 Visual Encoding Implementation

The relevance-based visualisation is shown in Figure 16. The relevant POIs are rendered with large markers featuring vibrant, theme-specific colours at complete opacity. These markers incorporate centred SF symbols that provide semantic identification through their colour hue. In contrast, irrelevant POIs are visualised through smaller markers and rendered in a semi-transparent grey, creating a clear visual recession that preserves spatial context without competing for visual attention.

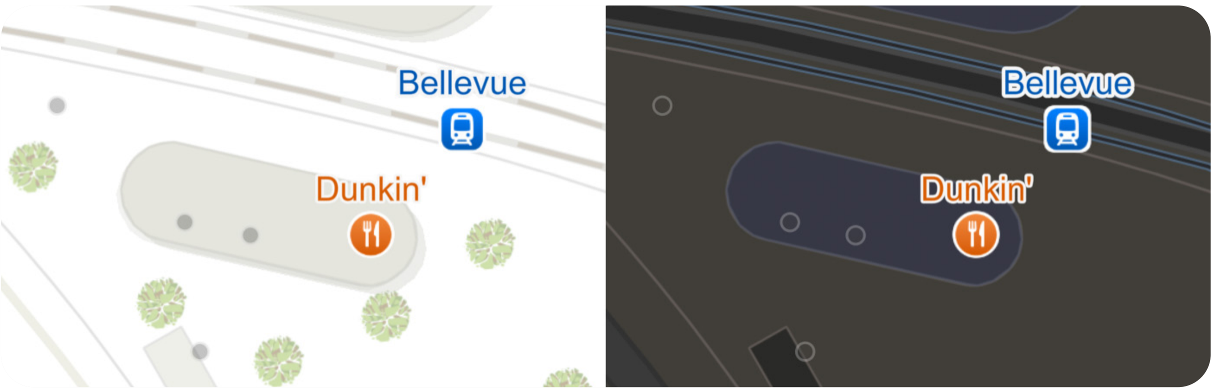

Figure 16: Close-up screenshot of the map visualisation with the relevant (tram stop and restaurant) and the irrelevant points of interest in day and night mode

The colour encoding reduces the complexity of 174 unique amenities into 12 semantically coherent groups, each assigned a distinct colour hue from a perceptually balanced palette inspired by the existing cultural associations (green for nature, blue for water or transportation) and colour choices from Apple Maps and OpenStreetMap. The implemented mapping assigns pink to cultural and leisure features, red to medical institutions, brown to educational POIs, green to places that are related to nature or outdoors, a light blue hue to sports, turquoise to tourism-related features, dark violet to religious institutions, yellow to shops, orange to restaurants and light violet to accommodations. The transportation POIs receive a blue colour hue and feature a distinct rectangle marker, while general infrastructure POIs are coloured in blue or grey. The POI symbology, including these colour assignments and marker designs, was created in Adobe Photoshop. An example of the amenity mapping can be seen in Figure 17. The exact mappings are in Appendix 4.

Figure 17: Overview of the point of interest amenities and the respective colour hue and an example symbology

4.1.3 Interactive Elements

The application's interactive design implements gesture-based touch controls and responsive feedback mechanisms that align with the ones that align with the iOS platform conventions while supporting the specialised needs of the adaptive map interaction. The map itself can be manipulated with simple pinch-to-zoom gestures or double taps that zoom in or out of the map. Pan navigation updates the visible extent while maintaining the POI selection states.

The geocoding integration makes use of the ArcGIS World GeocodeServer[7] to provide real-time location search capabilities. The search element uses several affordances to communicate functionality. A magnifying glass icon serves as a semantic signifier, while the placeholder text provides instructional context. As soon as text is entered, a blue search text appears and allows for a clear action availability.

[7] https://developers.arcgis.com/swift/geocode-and-search/, accessed 25.09.2025.

Successful geocoding triggers a coordinated response sequence and results in a blue circular marker appearing at the searched location with the corresponding text label positioned above (Fig. 18). The scale gets slightly smaller compared to the application startup to provide a better geographic orientation. A return from the search state and virtual location to the real-time user position happens through the click of the dedicated blue location button at the top right corner, as shown in Figure 18.

During the computationally intensive relevance calculation, a modal overlay informs the user about the current process with an activity indicator and accompanied text (Fig. 19). The overlay maintains map visibility but temporarily enforces an interaction constraint, preventing user actions that might conflict with the calculation pipeline.

The day/night toggle executes basemap transformations while maintaining the spatial arrangement and extent of the current map view. The visual feedback is supported by the icon change (sun/moon), which clearly indicates the current mode and supports manual correction according to different lighting conditions.

The theme picker implements a native iOS menu component that displays the six thematic options through a standard selection interface (Fig. 20). The current theme appears as text within the pill shaped selector, with a chevron indicator signifying interaction availability. After a selection, the manual thematic choice gets saved and a new adaption process is triggered.

Figure 18: Search result view of the user interface

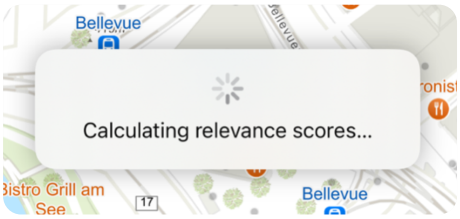

Figure 19: Loading overlay during computationally intense calculations

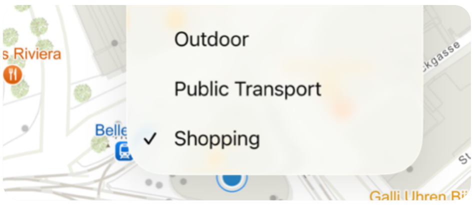

Figure 20: Expanded theme picker with the current theme marked with a check mark

4.1.4 Adaptive Behaviour

The application's adaptive capabilities are designed to respond dynamically to changing user contexts, with the implementation demonstrating theoretical variations in POI selection based on temporal, behavioural, and spatial factors.

The implemented adaptation mechanism produces theoretically predictable variations in POI selection and presentation based on the modelled user context. For example, when the system operates in an urban environment during weekday lunch hours, the thematic model predicts 'food' with high confidence, resulting in restaurants and cafes receiving semantic relevance boosts. At 09:00 in the same location, the TM shifts to prioritise 'public transport', with the RM predicting a high relevance score for all the transit stops. Opening the application on Friday evening, the cultural activity is predicted. Examples were virtually modelled with real predictions for the location of Zurich Bellevue, as shown in Figure 21.

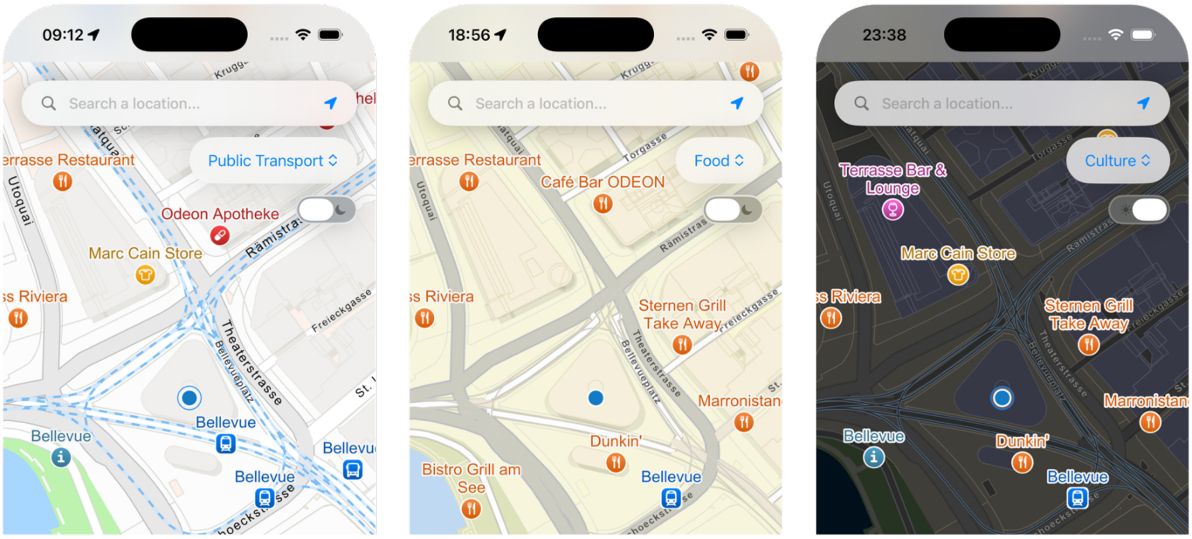

Figure 21: MyPlaces with different predicted thematic interests throughout the day

The behavioural adaptation mechanism is implemented through the relevance scoring pipeline, where user interactions modify POI scores according to the CPL framework from Section 3.3. The system design specifies that a restaurant with an initial relevance score of, for example, 0.5 would increase to approximately 0.7 after user click interactions (based on the click rate calculation) and further to around 0.85 when marked as favourite (through the behavioural component maximum selection). On the left side of Figure 22, the relevance was calculated without and with having the Carhartt store selected as a favourite.

Additionally, having two different users initialised that have different places saved to their favourites and feature different interaction histories, they receive a slightly adjusted map view despite the rest of the context variables being the same (location, time, theme). This can be seen on the right two images of Figure 22.

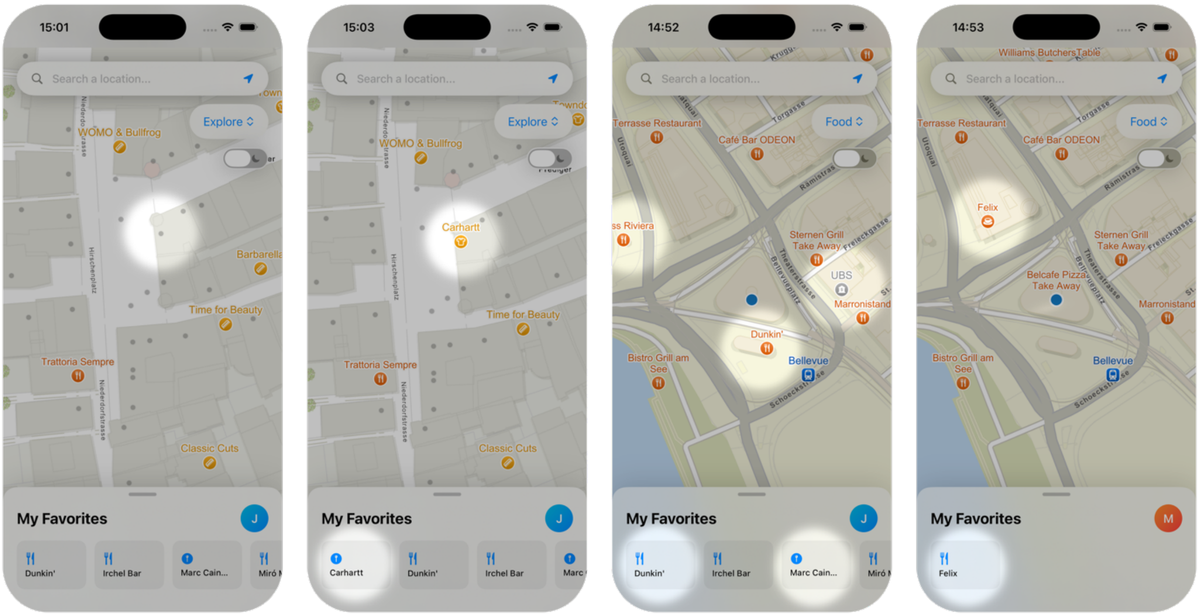

Figure 22: Influence of behavioural attributes with changes in click counts and favourites with the same user (left two screenshots) and different users (right two screenshots)

The aggregation behaviour adapts to both the map scale and POI density, producing distinct POI patterns across zoom levels. At street level (1:500), the system displays individual features with minimal aggregation. As the scale decreases to the neighbourhood level, aggressive clustering reduces the displayed POIs, prioritising highest-relevance features within each spatial cluster. Looking at Figure 23 with the user location set to the densely populated Zurich Bellevue, the consistent visual clarity can be observed across three different zoom levels.

The scale-dependent aggregation rendering hides the irrelevant POIs at scales smaller than 1:1’200 and only lets them appear when the user zooms in sufficiently to indicate detailed exploration intent. The relevant POIs also receive a label in a matching colour hue when scales larger than 1:1’500 are reached.

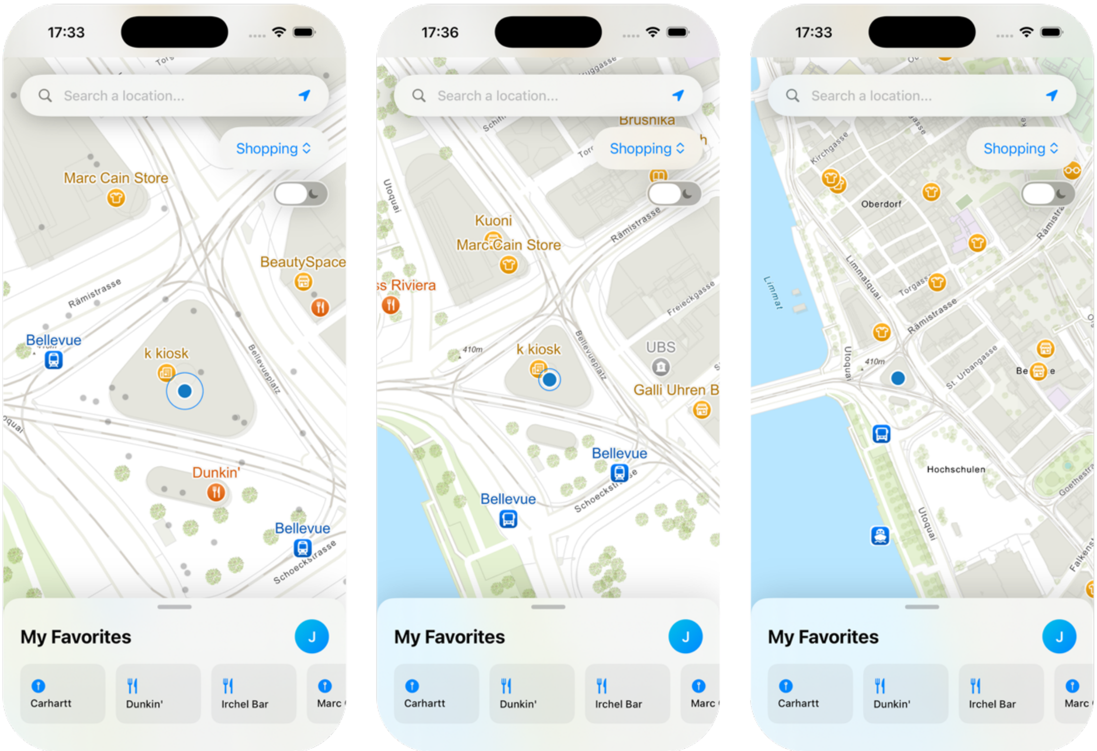

Figure 23: Three different scales that showcase the aggregation and relevance mechanism
Starting with a very small scale, small scale and medium scale (from left to right)

It has to be considered that all the screenshots were taken under partly simulated contexts, with the location and user profiles being virtually set to qualitatively assess the application behaviour and context adaptation. The displayed variations reflect the intended system behaviour as designed through the rule-based system and model training.

4.1.5 Scientific Contributions

The application implements support for further research through multiple mechanisms. First, the code is publicly available so that the implementation is replicable. It can be retrieved here: https://gitlab.uzh.ch/giva/public/masters/mobile_map_adaptivity_app_johannes_guler
The code itself is documented, and concise comments should help understand or modify specific functionality. The object-orientated code makes it easier to exchange the ML models, implement new features, or experiment with the relevance algorithms.
The SQLite database persistence layer captures complete interaction histories with all the relevance scores, click counts, the last clicked date of the POI and all the contextual attributes

used by the ML model for the prediction. A sample of the core data extract can be seen in Appendix 5.

During operation through the Xcode Simulator, the application provides real-time visibility into system behaviour. Each step of modelling, cache update or aggregation is printed to the console with insights into the numeric variables to provide a quick overview of the current processes.

4.2 Relevance Model Training & Output

This section presents the results from training and evaluating the ML models that power the adaptive functionality of MyPlaces. Both the TM and RM were trained and evaluated as outlined in Section 3.6.1.

4.2.1 Thematic Model Accuracy

The TM was trained to classify user activities into six thematic categories based on three integer-encoded input features: physical environment type (0–2), time of day (0–23), and day of the week (0–6). The synthetic training dataset includes 1'000 samples with augmentation expanding to 4'000 samples through noise injection.

4.2.1.1 Thematic Cross Validation Performance

Initial cross-validation using five-fold stratified sampling achieved a mean accuracy of 65.7%, with individual fold accuracies ranging from 63.6% to 67.5%. The 95% confidence interval is between [62.9%, 68.9%]. The grid search optimisation across 432 parameter combinations identified optimal hyperparameters: 200 estimators, a maximum depth of 3, a learning rate of 0.15, and 60% sub-sampling with 80% column sampling per tree. The overfitting analysis revealed a training accuracy of 69.3% in comparison to the validation accuracy of 66.6%, yielding a gap of 2.7%.

4.2.1.2 Class specific performance

The final model achieved 68.8% validation accuracy with varying performance across the six thematic categories. As shown in Table 4, the model had the strongest performance for the 'outdoor' (F1 = 0.786) and 'explore' (F1 = 0.800) themes, while 'food' classification proved the most challenging to predict (F1 = 0.353).

Table 4: Class performance overview of the thematic model training

Class	Precision	Recall	F1-Score	Support
shopping	0.568	0.462	0.509	91
food	0.517	0.268	0.353	56
public transport	0.645	0.652	0.649	184
culture	0.559	0.567	0.563	67
outdoor	0.73	0.851	0.786	194
explore	0.783	0.817	0.8	208
accuracy		0.688		800

4.2.1.3 *Thematic Feature Importance*

The feature importance in Figure 24 indicates that the physical environment type contributed 52.1% of the predictive power, time of day contributed 34.1%, and day of week contributed 13.8%.

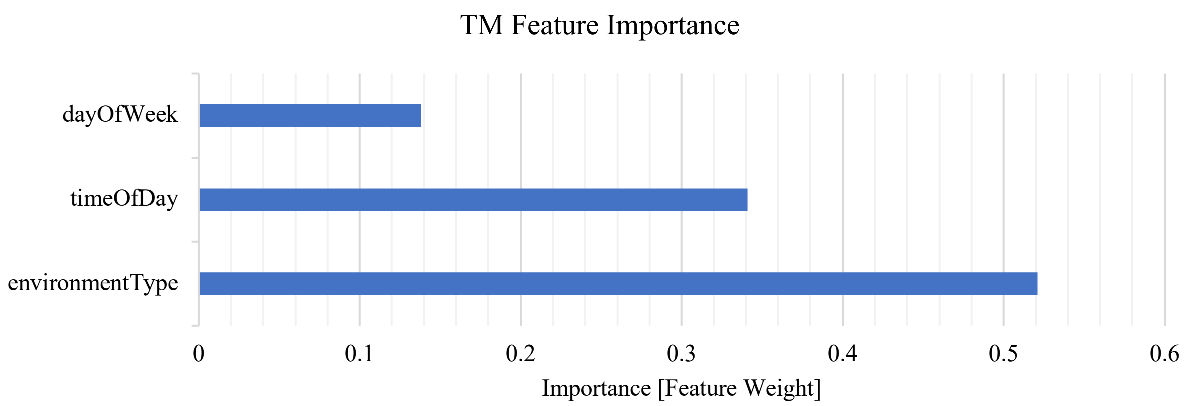

Figure 24: Feature weight importance of the thematic model

4.2.1.4 *Noise Robustness of thematic Predictions*

The TM maintained a robust performance under input noise conditions (Fig. 25). At baseline (0% noise), both rule-based and ML approaches achieved approximately 90% accuracy, with the rule-based prediction scoring slightly higher. However, as soon as noise increased up to 30%, the ML model degraded to 61% accuracy, while the rule-based prediction accuracy was even lower at 57%.

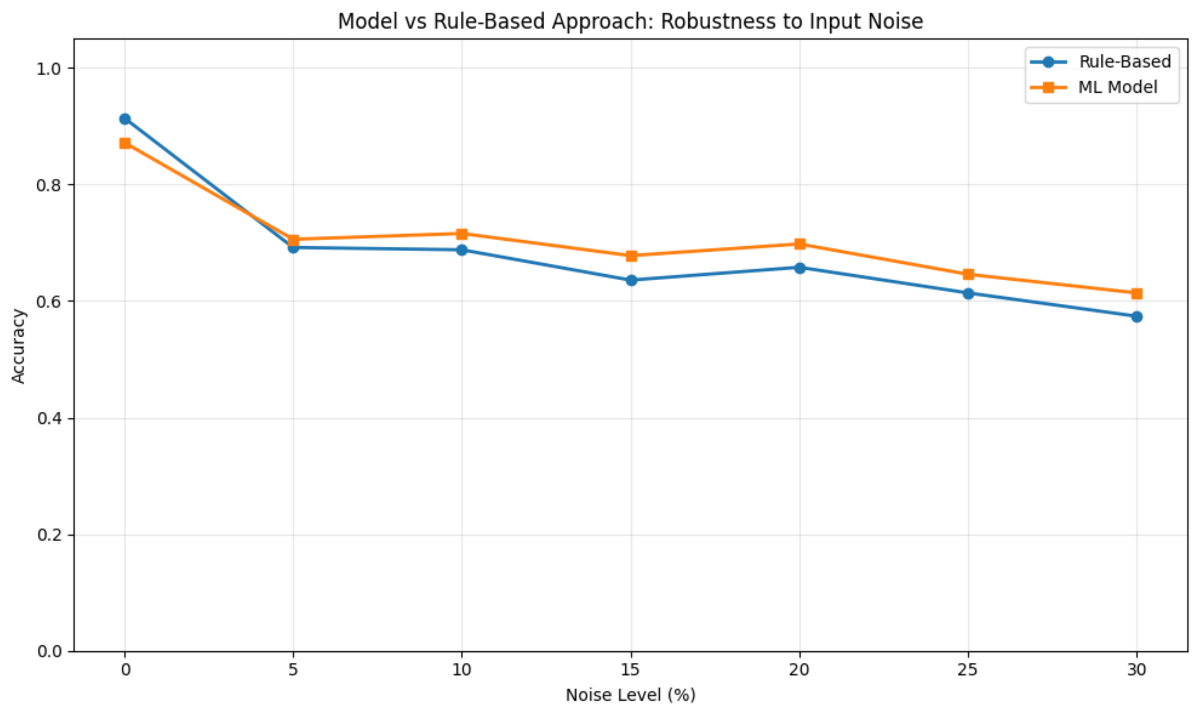

Figure 25: Model vs. rule-based thematic predictions with different input noise

4.2.2 Relevance Model Accuracy

The RM processed 11 input features to predict a continuous relevance score, ranging from 0 to 1. Training utilised 8’000 synthetic samples augmented to 16’000 through controlled noise injection.

4.2.2.1 Pre-Computed Geographic Feature Attributes

Prior to relevance modelling, cluster and co-location scores were computed for the entire POI dataset. Cluster score statistics revealed a mean score of 0.779 with a standard deviation of 0.341. Nearly all of the POIs received a non-zero cluster score (99.98%). Figure 26 reveals that 73.8% of the POIs received high scores above 0.8.

Co-location scores were overall lower with a really low mean of 0.009, which means that only around 2% of the POIs demonstrated significant co-location patterns.
The co-location rule mining identified 314 high-quality spatial association patterns from an initial candidate set of 10’802. Top scoring patterns included tourist information centres and viewpoints as well as hotels and restaurants due to complementary amenity proximity.

Table 5: Feature attribute metrics of cluster and co-location scores

Metric	*Cluster Value*	*Co-Location Value*
Mean	0.779	0.009
Std Dev	0.341	0.066
Min	0	0
Max	1	1
Median	0.961	0
Non-zero	1292938 / 1293088	24572 / 1293088

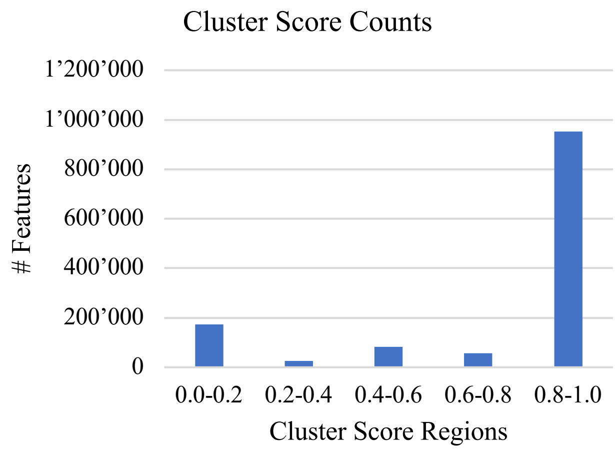

Figure 26: Cluster score count distribution across scoring regions

4.2.2.2 Relevance Model Training Metrics

Cross-validation metrics showed a mean RMSE of 1.345 (SD = 0.0053) across five folds. The hyperparameter optimisation evaluated 2,187 candidates through 6,561 model fits. The overfitting check revealed a training RMSE of 0.0759 and a validation RMSE of 0.1305, with a gap of 0.0546. The final model achieved strong regression performance with an RMSE of 0.0874, an MAE of 0.0531 and an R^2 of 0.6954. Prediction statistics demonstrated appropriate range coverage with the actual scores between [0.000, 0.866].

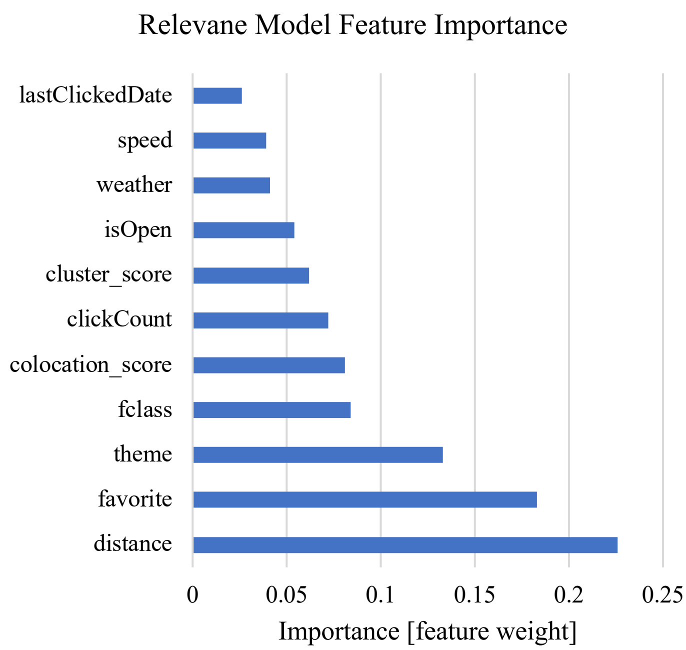

Figure 27: Feature importance of the relevance model

Feature importance analysis in Figure 27 reveals that distance is the most dominant factor (22.6%), followed by the favourite status (18.3%) and the predicted theme of the TM (13.3%). Geographic environment features (cluster, co-location) in combination contributed 14.3% to the modelled relevance. The least important is the last clicked date of the POI, with only 2.6% feature weight importance.

4.2.2.3 Computational Performance of the Relevance Model

Prediction speed measurements (Fig. 28) demonstrated a performance trade-off between approaches. Single predictions were both fast, with 0.27 ms for the ML and 0.02 ms for the rule-based function. At scale 1:500, the gap opened, and the ML needed 97.25 ms while the rule-based function only needed 2.20 ms. Finally, 5,000 predictions required 934.71 ms versus 21.20 ms, respectively. Overall, both models remained below a second of prediction time.

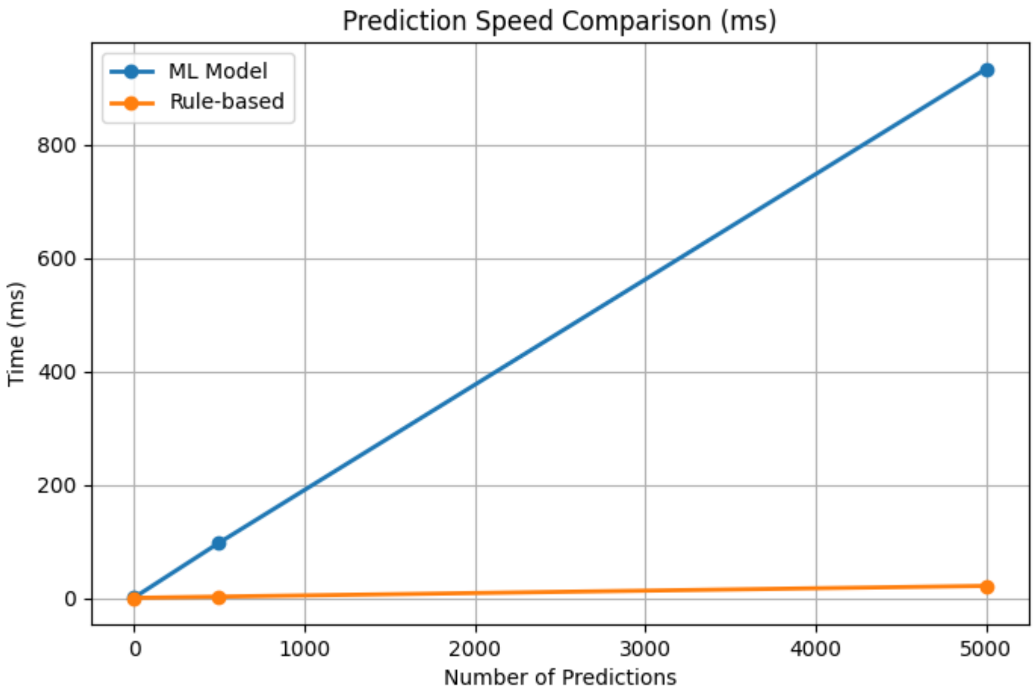

Figure 28: Prediction speed measurements for the machine learning model and rule-based function

Under noise injection testing (Fig. 29), both models maintained high-performance scores at baseline predictions. The performance was measured with a bound MSE. The ML model showed marginally superior stability across all noise levels up to 30%.

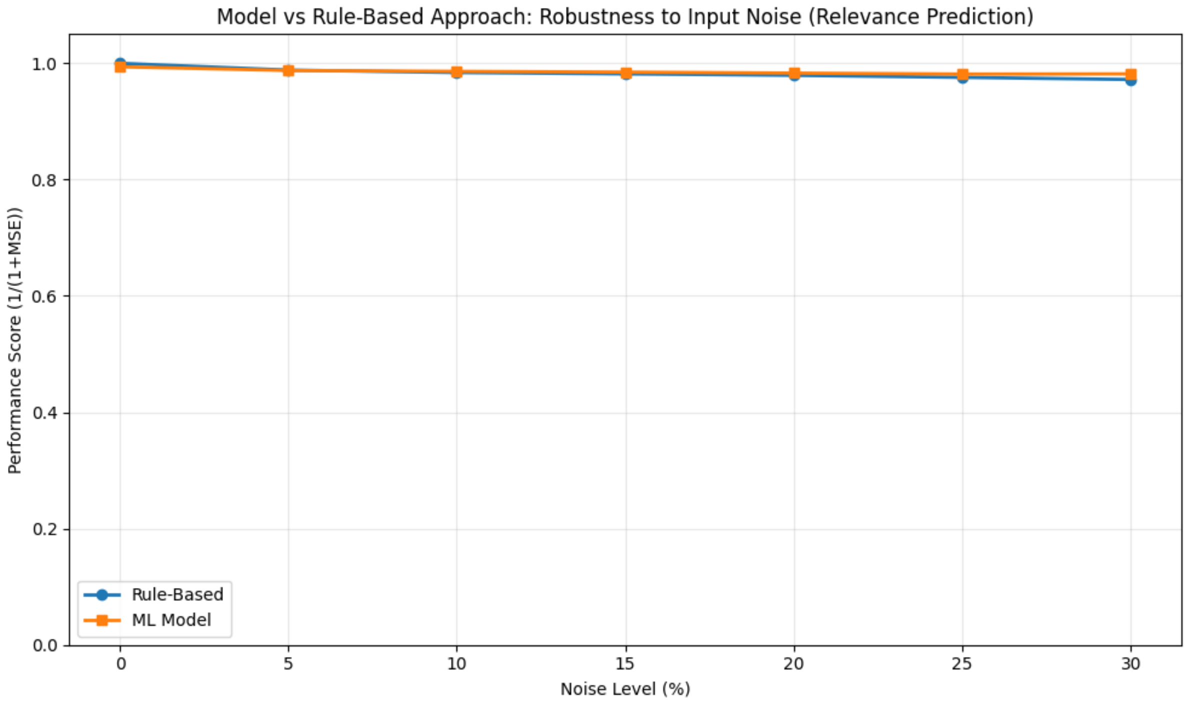

Figure 29: Model vs. rule-based relevance predictions with different input noise

4.3 Computational Performance Analysis

The computational performance analysis revealed application adaptation and aggregation patterns in different environments and scales, demonstrating both the strengths and limitations of the ML model approach. The whole dataset of the performance analysis can be found in Appendix 6. It was conducted inside Xcode and the corresponding simulator application of an iPhone 17 Pro with a fast and stable internet connection.

The scale influence analysis shown in Figure 30 granted insights into the POI loading and aggregation in an example of an urban environment. The same location of Zurich Bellevue was chosen analogous to Figure 23 during the qualitative aggregation evaluation.

At a scale of 1:500, 425 POIs were loaded onto the map, reducing the totally available POIs by 73%. Zooming further out to a scale of 1:1’500, the application visualised 318 displayed POIs (80% reduction). This number was further reduced at a scale of 1:5’000 with a reduction of 96% of the available POIs and an absolute amount of 66 POIs.

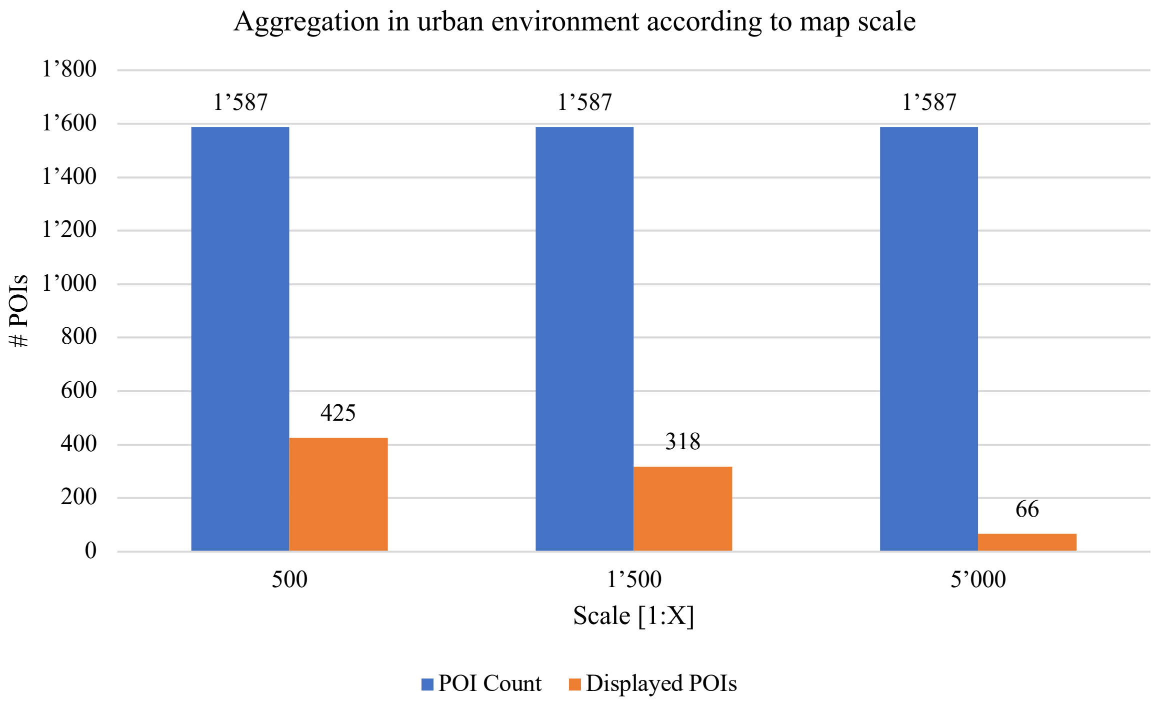

Figure 30: Aggregation of points of interests according to map scale at Zurich Bellevue

A detailed breakdown of the time distribution analysis across four main components is presented in Figure 31. The split is made according to the same urban location in Zurich for a scale of 1:500. The aggregation component dominates the processing pipeline, consuming approximately 5.1 seconds (56% of the processing time). This component analyses the spatial relationships between all loaded POIs based on relevance scores and reduces the visualised POIs from 1,587 to 859. The second most time-consuming component is the relevance prediction phase at 1.5 seconds (16% of total time). This represents the ML model inference for all 1,587 POIs, calculating personalised relevance scores based on distance, user preferences, temporal factors and the geographic environment patterns. These attributes are calculated on average in 0.93 ms per POI. The POI loading from ArcGIS Location Platform consumes roughly 1.4 seconds and represents 15% of the total time. This is mainly influenced by the network latency and data transfer for fetching POI attributes within the POI buffer zone. The thematic prediction requires 1.1 seconds (12% of the total time) for the ML model inference.

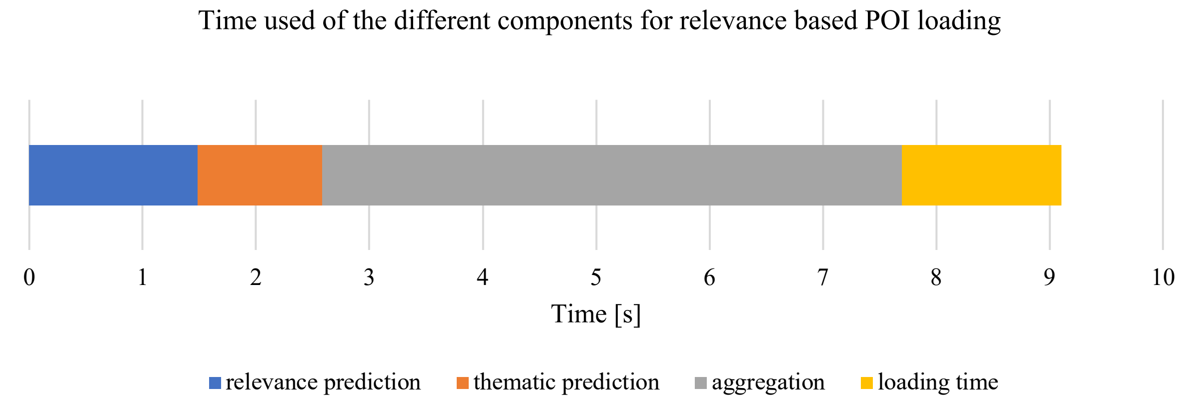

Figure 31: Different contributors of the geographic relevance assessment of points of interests with an urban location

As visualised in Figure 32, urban environments present the most computationally demanding scenarios with a consistently high number of POIs across all zoom levels. Total processing time ranges from 7.7 to 9.1 seconds, with the highest processing time at the most detailed scale (1:500). The computation time decreases as the scale becomes coarser, reaching 7.7 seconds at 1:1,500 and 8.7 seconds at 1:5,000. Rural settings improve the POI computations drastically with only 116 POIs to process. The time varies between 0.8 and 1.3 seconds, showing excellent performance across all scales. Outdoor environments achieve near-instantaneous performance with only 9 POIs in the loading buffer. Processing times remain minimal across all scales, ranging from 0.3 to 0.7 seconds, with the computation time remaining consistently below one second regardless of scale changes.

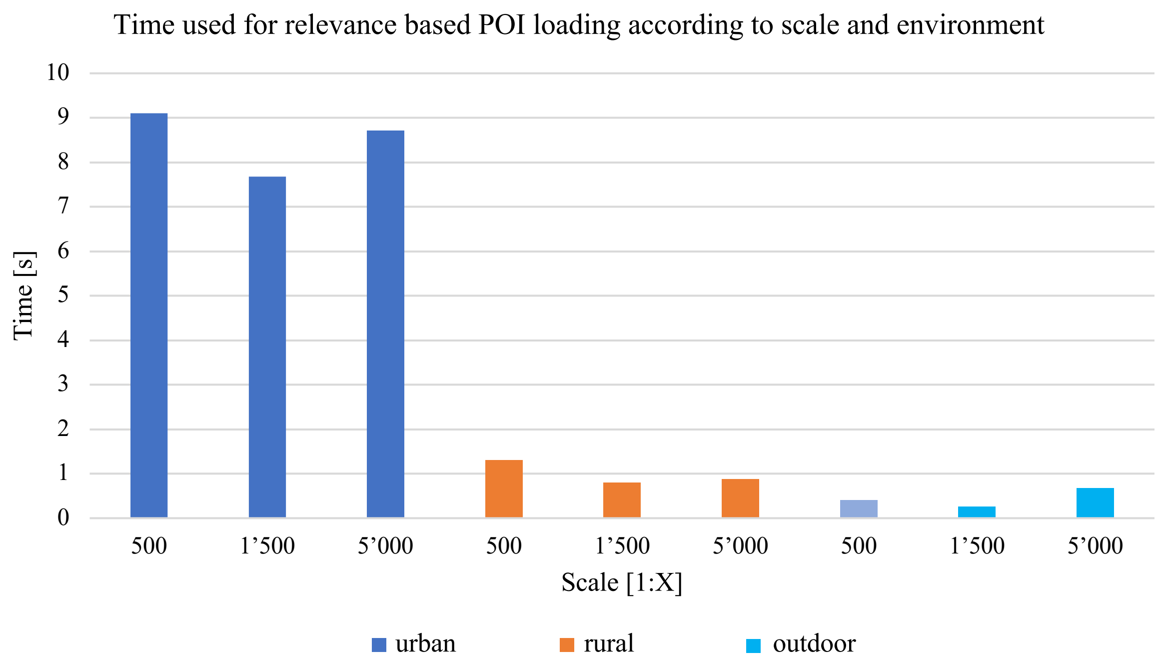

Figure 32: Loading times of points of interests with the geographic relevance calculation in urban, rural and outdoor environments with different map scales

5 Discussion

This chapter critically analyses the outcomes of the MyPlaces prototype implementation, evaluating its ability to address the research questions and demonstrating the feasibility of real-time adaptive mapping. The discussion combines the computational and practical results with the theoretical framework to position MyPlaces as both a proof-of-concept and a foundation for future research.

5.1 Methodological Reflections

5.1.1 Geographic Relevance Modelling

The implementation successfully demonstrates that GR can be effectively modelled through a multi-dimensional approach that combines the different dimensions of user context. The adopted framework of the computational model from De Sabbata (2013) proved to be theoretically grounded and, with slight modifications, practically implementable. Relevance predictions received an R^2 of almost 0.7, despite being trained on synthetic data.

The theory-driven design was able to provide several insights into context modelling for the GR assessment. First, the CPL framework successfully integrated the mandatory context variables (spatio-temporal, semantic, behavioural) and the desired factors (geographic environment), producing relevance scores that aligned with the theoretical expectations.

The feature importance analysis validates this approach: distance dominates the GR prediction, followed by user interaction indicators (favourite status) and thematic alignment. This distribution profoundly validates Tobler's (1970) First Law of Geography, demonstrating that even within a multidimensional relevance framework, spatial proximity remains the primary driver. The implementation also reveals departures from the theoretical expectations. While De Sabbata and Reichenbacher (2012) identified co-location as one of the main GR factors, only a fraction of the identified co-location rules demonstrated meaningful, resulting in a minimal impact on the relevance score. Combined with the uneven distribution of cluster scores, these outcomes suggest an influence of the OSM POI data characteristics, which naturally clusters around densely inhabited and frequently visited locations. This indicates that both data quality and properties must be considered when computing GR.

The approach to Bartling et al.'s (2023) context taxonomy reveals both achievements and limitations. The system successfully captures various extrinsic factors (location, time, movement speed, weather) and behavioural interactions (clicks, favourites, searches). The behavioural factor is assessed with the activity prediction of the TM to bridge the missing direct sensing and assess it with derived extrinsic attributes. Despite the 85% accuracy, the synthetic data generation inherently limited the model's ability to learn genuine user preference patterns. This is also reflected in the relatively high RMSE of the RM, which suggests room for improvement with real-world training data. As Fabrikant (2023) argues and Zingaro et al. (2024) empirically demonstrate, individual differences in spatial abilities, cognitive load, and spatial anxiety fundamentally shape map interaction patterns. The current implementation cannot distinguish between a stressed user needing simplified information and a relaxed user seeking comprehensive options.

5.1.2 Technical Evaluation

The technical architecture with the combination of Swift and the ArcGIS Maps SDK proved well suited for implementing adaptive mobile mapping, though with identifiable bottlenecks. With the concept of OOP, concerns between context sensing, relevance computation, and visualisation are successfully separated. This enables modular experimentation possibilities for further research.

The ML learning architecture marks a shift from the previous rule-based approaches and demonstrates, that sophisticated GR assessment can operate within mobile constraints. The performance analysis revealed that ML computation generates a computational overhead, yet the absolute times remain acceptable: The lower loading times of 0.8-1.3s in rural and 300-700ms in outdoor environments maintain what Attig et al. (2017) identify as crucial for task success. However, urban environments with almost 2'000 POIs and a loading time of just below 10 seconds push beyond comfortable boundaries. Most of that time is however generated through the aggregation process with the $O(n^2)$ spatial evaluations. This suggests that future optimisations should prioritise an efficient spatial aggregation over model efficiency.

The decision to employ ML models over pure rule-based systems proves reasonable, despite the performance penalty. The models offer two critical advantages that outweigh this cost. First, the capacity for continuous learning from user behaviour enables evolution beyond designer assumptions. The XGBoost model can capture relationships beyond the static rules and discover new patterns in user-context-relevance relationships. Second, the models' ability to handle non-linear interactions and missing data proves essential when working with data sources like OSM, where attribute completeness varies dramatically.

5.2 Application Implications

5.2.1 Practical Implications

The MyPlaces prototype successfully demonstrates dynamic adaptation through coordinated visual encoding and real-time relevance computation of POIs. The visual hierarchy creates perceptual distinctions between relevant and irrelevant places through variations in size, opacity, and colour saturation. This aligns with Bertin's (1983) "visual hierarchy of importance" and with Garlandini and Fabrikant's (2009) findings on size effectiveness for mobile displays. This relevance encoding of the POI establish a figure-ground distinctions in line with Gestalt theory.

The aggregation algorithm extends cartographic generalisation theory by incorporating relevance into reduction decisions, creating "relevance-aware generalisation." Through a change in scale, the app reacts according to the viewing context. Interestingly, the aggregation time increases slightly at larger scales in urban environments, which is likely due to more complex clustering decisions across larger spatial extents.

The implementation fulfils all functional requirements (F1–F9) established in Section 3.2, though with varying degrees of sophistication. Context awareness (F1) and relevance computation (F3) with the RM, TM, and variable manager represent the most mature components, while search functionality (F7), POI popups (F8), and orientation guides (F9) remain basic but functional. The dual POI layer methodology and different basemaps provides the assessment of F4, where the night theme specifically addresses the luminosity adaption pointed out by Zhang et al.'s (2009) research. The implemented feedback mechanisms, such as toggles, loading overlays, or the theme picker, assess F5.

Looking at the technical requirements, object-orientated design principles (T1) enable the promised modularity, with clean separation between data, logic, and presentation layers. Processing all the relevance assessments on-device guarantees privacy preservation (T5) and addresses the highlighted growing user concerns by Ul Haque et al. (2023) about data misuse and location tracking. Performance optimisation (T3) with caching (T2) achieves acceptable real-time behaviour, though urban environments with large potentially relevant places push the boundaries of instant response.

The three defined use cases demonstrate varying levels of success:

- Use Case 1: On-the-go relevance discovery performs well in practice. The system identifies and highlights POIs considered as relevant in less than 10 seconds in feature-rich environments. While exceeding an ‘instant’ response, this latency is likely acceptable for exploratory tasks where users probably expect certain processing times.
- Use Case 2: Revisiting favourite places shows good performance. The favourite weighting ensures previously saved locations maintain high visibility and are likely to be highlighted as relevant POIs across different contexts. Additionally, the implementation in the Favourite Panel guarantees constant access.
- Use Case 3: Location exploration and planning reveal limitations in the current implementation. Despite a functioning search function, the GR assessment is made with the same models as with the real user location, resulting in less accurate predictive results.

5.2.2 Scientific Implications

The MyPlaces prototype is a feasible research platform for investigating adaptive cartography and GR. Several architectural decisions specifically support future research:

1. A modular ML pipeline separates training scripts and model integration to enable a quick model exchange without application recompilation. Researchers can experiment with alternative algorithms, feature sets, or training strategies.
2. Comprehensive data logging to the SQLite database captures user-relevance-POI triplets, building a dataset to gain insights into real user behaviour.
3. The performance collector provides a detailed timing breakdown for optimisation research. With the benchmark tests, future improvements can be conducted.
4. While currently limited to extrinsic factors, the context modelling is extendable and allows for additional sensors and intrinsic factors to be included.

MyPlaces provides quantifiable relevance metrics rather than hidden decisions so that application responses can be retraced and understood. Multi-user profile support enables comparative studies and analysis of behavioural differences on the modelled GR.
The scientific requirements (S1–S4) are comprehensively met. Experimental validity (S1) emerges through detailed logging of all relevance computations with the context variables. Reproducibility (S2) is ensured through version control and comprehensive documentation. The modular architecture (S4) already proved its value during development, with the relevance engine undergoing major changes without affecting other components.

5.3 Research Contributions and Limitations

5.3.1 Limitations in Context Modelling

The model training is exclusively based on theoretically based synthetic data. This creates a ceiling effect, where the model can, at best, match the rules that created its training data. More realistic modelling should be aimed at, with noise injection. Nevertheless, the high baseline performance (R^2 = 0.695) may not translate to real-world scenarios where user behaviour deviates from synthetic patterns.
Additionally, one of the key limitations already mentioned in 5.1.1 is the lack of intrinsic and behavioural factors for the context modelling. The prototype insufficiently captures internal user states such as cognitive load or user activity. The latter is partially assessed through the TM, but without psychological sensors or self-reported state indicators, the system cannot adapt to user stress or confusion. Fabrikant's (2023) vision for neuroadaptive maps and Zingaro et al.'s (2024) empirical findings on the importance of individual spatial abilities highlight the criticality of these missing dimensions.
The performance measurements assumed a fast and stable internet connection, but real-world scenarios face latency and variability in internet speed. The loading time could extend several seconds, breaking the responsiveness requirement.

5.3.2 Limitations in Application Design

The main limitation of the application design is the lack of formal user testing. The assumed benefits of adaptation and increase of usefulness in the application remain theoretical. Users might find automatic POI filtering disorienting or prefer explicit control over map content.

While successfully implementing two of Reichenbacher and Bartling's (2023) three adaptation dimensions (geographic information and visualisation) the prototype is limited in interface adaptation. The basemap adjusts to the activity only in certain implemented scenarios (outdoor, public transport) and the icon colours and sizes do not account for user visual acuity. The UI layout also does not adapt to the activity, and the elements remain in their hard-coded form and position.

5.3.3 General Limitations

As previous adaptive applications have struggled with user trust and transparency, the automatic adaptation of the visual hierarchy of POIs could generate confusion. Users cannot see why certain POIs appear or disappear, and the black box of ML predictions compounds this issue. While data remains on the device, users might still feel uncomfortable with detailed behaviour tracking.
In the current implementation, there is minimal accessibility support. There are no voice guidance or contrast enhancement options, and the text sizes are fixed. This excludes users with visual impairments who could most benefit from intelligent spatial information filtering.

5.4 Outlook

MyPlaces as a prototype establishes a foundation for multiple research directions. Context modelling can be enhanced with the integration of additional intrinsic sensing techniques, such as wearable sensors (e.g., heart rate for stress) or questionnaires, which could dramatically improve relevance assessment. The dimension of directionality, currently absent, could be implemented as an additional factor to predict places along the movement path as soon as the user is not stationary. To make the application more responsive, additional performance improvements, such as batch loading or progressive POI loading could be implemented. Deploying the application to gather authentic user interactions would enable training models that capture real-world behaviour patterns. A feedback loop where manual theme changes train the TM would create continuous improvements.

Further versions could adapt the entire interface and adjust the basemap more specifically to the modelled thematic user interest and cognitive load. High-contrast modes or adjustable simplification could make the application accessible to an even larger audience. The user profile settings could be used, where users can modify their own preferences.

Large language models could enable natural language queries about spatial information or could be used to suggest potentially relevant places in the current and future situations, while generative AI models might create personalised map styles that match user preferences.

6 Conclusion

This thesis aims to demonstrate that adaptive mobile maps based on GR are not only theoretically justified but also technically feasible in real-time mobile environments. With the help of MyPlaces, a functional AMMA iOS prototype, this goal has been achieved.

The application implements real-time adaptive mapping through an object-orientated architecture which processes user context data to assess POI relevance. Both map content and visual encoding are dynamically adjusted. The visual adaptation strategy uses size, opacity, and colour saturation as variables to create clear relevance hierarchies, while an aggregation algorithm reduces visual complexity. Multi-user support and data logging for further analysis are additional features that are implemented. To preserve privacy, relevance assessments are carried out on-device.

The implementation successfully integrated and adjusted the proposed GR modelling framework from De Sabbata (2013). It uses CPL to combine both mandatory and desired context attributes, of which distance, user preferences, and topicality are the most influential factors. The event-driven architecture achieves dynamic adaptations that are based on real-time relevance assessments with two trained ML models. The hosted geographic data are filtered and visualised to provide an adjusted mobile map view.

With the modular design, ML integration and comprehensive data logging, the prototype application can be used to gain context modelling insights. For the GIVA research group at the University of Zurich, MyPlaces provides a concrete prototype tool to further investigate AMMAs. The open-source codebase ensures reproducibility and encourages community contribution.

Key limitations of the application are the heavy reliance on the synthetic training data and absence of user evaluation studies. The current implementation also mainly models the user context based on extrinsic context variables. Mobile network availability and limited accessibility features further constrain real-world deployment.

Future research should, therefore, prioritise two main directions. First, research studies could collect real-world interaction data to train the models on authentic user behaviour patterns. Second, the application adaptivity should be further expanded to the UI elements (text and icon size, basemap styles) and include accessibility features, such as voice guidance or colour contrast boost.

Looking ahead, the transition from static to adaptive maps represents a fundamental shift in information interaction. As location-based services become increasingly important in daily life, the need for intelligent information filtering grows. MyPlaces demonstrates that GR provides a functional approach to this challenge, balancing information richness with cognitive manageability. The future of mobile mapping lies in revealing the information that matters most, when it matters most. This thesis provides a stepping stone towards that future. What remains is the research, development, and refinement needed to transform adaptive mapping from a promising prototype to a ubiquitous reality.

References

Abowd, G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., & Steggles, P. (1999). Towards a better understanding of context and context-awareness. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *1707*, 304–307. https://doi.org/10.1007/3-540-48157-5_29

Attig, C., Rauh, N., Franke, T., & Krems, J. F. (2017). System latency guidelines then and now: Is zero latency really considered necessary? *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *10276 LNAI*, 3–14. https://doi.org/10.1007/978-3-319-58475-1_1/TABLES/1

Atzmanstorfer, K., Eitzinger, A., Marin, B. E., Parra Arteaga, A., Quintero, B. G., & Resl, R. (2016). HCI-evaluation of the GeoCitizen-reporting app for citizen participation in spatial planning and community management among members of marginalized communities in Cali, Colombia. *GI_Forum*, *1*. https://doi.org/10.1553/giscience2016_01_s117

Awasthi, G. K., Tewari, S., S, N., P, G., Mallick, S., & S, S. (2025). Human-computer interaction design principles for enhancing user experience in mobile applications. *ITM Web of Conferences*, *76*, 05003. https://doi.org/10.1051/itmconf/20257605003

Ayada, W. M., & Ezz Eldin Hammad, M. A. (2023). Design quality criteria for smartphone applications interface and its impact on user experience and usability. *International Design Journal*, *13*(4), 339–354.

Ballatore, A., McArdle, G., Kelly, C., & Bertolotto, M. (2010, 22–26 March). *RecoMap: An interactive and adaptive map-based recommender* [Paper presentation]. Proceedings of the 2010 ACM Conference on Applied Computing (SAC), Sierre, Switzerland, 887-891. https://doi.org/10.1145/1774088.1774273

Bartling, M., Reichenbacher, T., & Fabrikant, S. I. (2023). Leveraging map use context for advancing cartography in the 21st century. *Proceedings of the ICA*, *5*(2), 1–7. https://doi.org/10.5194/ica-proc-5-2-2023

Bartling, M., Resch, B., Eitzinger, A., & Zurita-Arthos, L. (2019). A multi-national human-computer interaction evaluation of the public participatory GIS GeoCitizen. *GI_Forum 2019*, *1*, 18–39. https://doi.org/10.1553/giscience2019_01_s18

Bartling, M., Resch, B., Reichenbacher, T., Havas, C. R., Robinson, A. C., Fabrikant, S. I., & Blaschke, T. (2022). Adapting mobile map application designs to map use context: A review and call for action on potential future research themes. *Cartography and Geographic Information Science*, *49*(3), 237–251. https://doi.org/10.1080/15230406.2021.2015720

Bartling, M., Robinson, A. C., Resch, B., Eitzinger, A., & Atzmanstorfer, K. (2021). The role of user context in the design of mobile map applications. *Cartography and Geographic Information Science*, *48*(5), 432–448. https://doi.org/10.1080/15230406.2021.1933595

Beaudouin-Lafon, M. (1993). An overview of human-computer interaction. *Biochimie*, *75*(5), 321–329. https://doi.org/10.1016/0300-9084(93)90165-O

Beck, K., & Cunningham, W. (1989). A laboratory for teaching object-oriented thinking. *OOPSLA '89 Proceedings*, 1–6.

Been, K. (2002). *Responsive thinwire visualization of large geographic datasets* [Doctoral dissertation, New York University]. https://cs.nyu.edu/media/publications/been_kenneth.pdf

Bertin, J. (1983). *Semiology of graphics: Diagrams, networks, maps*. The University of Wisconsin Press.

Butler, C. (2023a). *Gestalt principles of design: Symmetry*. https://www.chrbutler.com/gestalt-principles-of-design-symmetry

Butler, C. (2023b, June 16). *Gestalt principles of design: Closure*. https://www.chrbutler.com/gestalt-principles-of-design-closure

Caro-Alvaro, S., García-López, E., García-Cabot, A., & Mavri, A. (2025). Measuring the effects of usability recommendations for mobile instant messaging apps on user performance and user behaviour. *International Journal of Human–Computer Interaction*, *41*(1), 608–626. https://doi.org/10.1080/10447318.2024.2303553

Chang, D., Dooley, L., & Tuovinen, J. E. (2002). *Gestalt theory in visual screen design: A new look at an old subject*. Proceedings of the 7th World Conference on Computers in Education: Australian Topics, 8. https://www.researchgate.net/publication/2532121_Gestalt_Theory_in_Visual_Screen_Design_-_A_New_Look_at_an_Old_Subject

Chen, B., Yu, S., Tang, J., He, M., & Zeng, Y. (2017). Using function approximation for personalized point-of-interest recommendation. *Expert Systems with Applications*, *79*, 225–235. https://doi.org/10.1016/J.ESWA.2017.01.037

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 13–17 August 2016*, 785–794. https://doi.org/10.1145/2939672.2939785

Clarke, K. C. (2001). Cartography in a mobile internet age. *The 20th International Cartographic Conference, ICC2001*, *3*, 1481–1488. http://www.geog.ucsb.edu/~kclarke

Cláudio, A. P., Carmo, M. B., Labmag, F. G., & Mendes Leal, M. (2011). Relevant points of interest: A visualization issue. *Revista de Ciências Da Computação*, *6*, 1–16.

Coalter, J. (2020). Colorbrewer 2.0 and the rainbow: Using color tools to choose appropriate color schema for your data visualization. *Issues in Science and Technology Librarianship*, *94*.

Couclelis, H., & Tobler, W. (1987). Exploring the anchor-point theory of spatial cognition. *Journal of Environmental Psychology*, *7*(2), 99–122. https://www.researchgate.net/publication/246147681

Crease, P. (2014). *Representation of geographic relevance in mobile applications* [Dissertation, University of Zurich]. http://opac.nebis.ch/ediss/20152471.pdf

Crease, P. E., & Reichenbacher, T. (2011). *Designing usable cartographic representations of geographic relevance for LBS users* (p. CO-073) [Paper presentation]. 25th International Cartographic Conference. http://icaci.org/documents/ICC_proceedings/ICC2011/Oral Presentations PDF/B1-Location based services and ubiquitous cartography/CO-073.pdf

De Sabbata, S. (2013). *Assessing geographic relevance for mobile information services* [Dissertation, University of Zurich]. https://www.geo.uzh.ch/dam/jcr:669c6ea0-dcb9-455f-817e-069224686a13/2013_Dissertation_DeSabbataS.pdf

De Sabbata, S., & Reichenbacher, T. (2012). Criteria of geographic relevance: An experimental study. *International Journal of Geographical Information Science*, *26*(8), 1495–1520. https://doi.org/10.1080/13658816.2011.639303

Durand, G. C., Ma, J., Pinnecke, M., & Saake, G. (2018). Piecing together large puzzles, efficiently: Towards scalable loading into graph database systems. *GI-Workshop on Foundations of Databases (Grundlagen von Daten- Banken)*. https://snap.stanford.edu

Elias, B., Hampe, M., & Sester, M. (2005). Adaptive visualisation of landmarks using an MRDB. *Map-Based Mobile Services: Theories, Methods and Implementations*, 73–86. https://doi.org/10.1007/3-540-26982-7_6

Fabrikant, S. I. (2023). Neuroadaptive mobile geographic information displays: an emerging cartographic research frontier. *International Journal of Cartography*, 1–17. https://doi.org/10.1080/23729333.2023.2253645

Forrest, D., & Castner, H. W. (1985). The design and perception of point symbols for tourist maps. *Cartographic Journal*, *22*(1), 11–19. https://doi.org/10.1179/caj.1985.22.1.11

Förster, M. (2023). *Mapping public transport stops in smartphone apps: Dynamic information visualization in the context of hiking*. [Master's thesis, University of Zurich.] https://www.zora.uzh.ch/entities/publication/a1f1f44e-7f5e-493b-9f6d-6592aa9b1aff

García, J. A. (2017). Time use patterns of Spanish people at weekends: In search of what, who and when. *Leisure Studies*, *36*(6), 793–810. https://doi.org/10.1080/02614367.2016.1252786

Garlandini, S., & Fabrikant, S. I. (2009). Evaluating the effectiveness and efficiency of visual variables for geographic information visualization. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *5756 LNCS*, 195–211. https://doi.org/10.1007/978-3-642-03832-7_12

Graham, L. (2008). Gestalt theory in interactive media design. *Journal of Humanities and Social Sciences*, *2*(1). https://www.guillaumegronier.com/2024-psychologiegenerale/resources/articles/Graham2008.pdf

Griffin, A. L., White, T., Fish, C., Tomio, B., Huang, H., Sluter, C. R., Bravo, J. V. M., Fabrikant, S. I., Bleisch, S., Yamada, M., & Picanço, P. (2017). Designing across map use contexts: A research agenda. *International Journal of Cartography*, *3*(sup1), 90–114. https://doi.org/10.1080/23729333.2017.1315988

Gyorödi, R., Zmaranda, D., Georgian, V., & Gyorödi, C. (2017). A comparative study between applications developed for Android and iOS. *International Journal of Advanced Computer Science and Applications*, *8*(11). https://doi.org/10.14569/ijacsa.2017.081123

Harrell, M. (2025, July). *From paper to AI: The evolution of maps and what's next*. Forbes. https://www.forbes.com/councils/forbestechcouncil/2025/07/07/from-paper-to-ai-the-evolution-of-maps-and-whats-next/

Harrie, L., Sarjakoski, T., & Lehto, L. (2002). A variable-scale map for small-display cartography. *International Archives of Photogrammetry Remote Sensing and Spatial Information Science*, *34*. https://www.researchgate.net/publication/228819722_A_variable-scale_map_for_small-display_cartography

Haswell, O., West, L., Savill, B., Whitham, R., & Martin, R. (2014). The effect of British Summer Time on outdoor physical activity. *The Lancet*, *384*, S35. https://doi.org/10.1016/S0140-6736(14)62161-9

Hatcher, G. S., & Mahmassani, H. S. (1992). Daily variability of route and trip scheduling decisions for the evening commute. *Transportation Research Record*, 72–81.

Hayir-Kanat, M., & Breuste, J. (2020). Outdoor recreation participation in Istanbul, Turkey: An investigation of frequency, length, travel time and activities. *Sustainability*, *12*(2). https://doi.org/10.3390/SU12020741

Hill, R., & Wesson, J. (2010). A-POInter: An adaptive mobile tourist guide. In *Proceedings of the 2010 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists (SAICSIT '10)*. New York: Association for Computing Machinery, 113–122. https://doi.org/10.1145/1899503.1899516

Hirsh, M., Prashker, J. N., & Ben-Akiva, M. (1986). Day-of-the-week models of shopping activity patterns. *Transportation Research Record*, 63–69.

Huang, L., Ma, Y., & Liu, Y. (2015). Point-of-interest recommendation in location-based social networks with personalized geo-social influence. *China Communications*, *12*(12), 21–31. https://doi.org/10.1109/CC.2015.7385525

Jenny, B., & Kelso, N. V. (2007). *61 cartographic perspectives Color Design for the Color Vision Impaired*. http://www.ryobi-sol.co.jp/visolve/en/

John, V. (2014). Functional role of colour in preattentive cognitive processing from thematic map images. In *IOSR Journal Of Humanities And Social Science (IOSR-JHSS* (Vol. 19, Issue 7). www.iosrjournals.org

Kaptelinin, V., & Nardi, B. (2007). Acting with technology: Activity theory and interaction design. *First Monday*. https://doi.org/10.5210/FM.V12I4.1772

Kiefer, P., Giannopoulos, I., Athanasios Anagnostopoulos, V., Schöning, J., & Raubal, M. (2017). Controllability matters: The user experience of adaptive maps. *GeoInformatica*, *21*(3), 619–641. https://doi.org/10.1007/s10707-016-0282-x

Konstantinou, E. N., Skopeliti, A., & Nakos, B. (2023). POI symbol design in web cartography: A comparative study. *ISPRS International Journal of Geo-Information*, *12*(7). https://doi.org/10.3390/ijgi12070254

Lewis, J., & Loftus, W. (2018). *Java software solutions: Foundations of program design* (9th ed.). Pearson International. www.pearsonglobaleditions.com

Li, H., Guensler, R., Ogle, J., & Wang, J. (2004). Using global positioning system data to understand day-to-day dynamics of morning commute behavior. *Transportation Research Record*, *1895*, 78–84. https://doi.org/10.3141/1895-11

Liu, B., Fu, Y., Yao, Z., & Xiong, H. (2013). Learning geographical preferences for point-of-interest recommendation. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, *Part F128815*, 1043–1051. https://doi.org/10.1145/2487575.2487673

Mackamul, E., Casiez, G., & Malacria, S. (2024). Clarifying and differentiating discoverability. *Human-Computer Interaction*, 1–16. https://doi.org/10.1080/07370024.2024.2364606

Meng, L. (2005). Ego centres of mobile users and egocentric map design. *Map-Based Mobile Services: Theories, Methods and Implementations*, 87–105. https://doi.org/10.1007/3-540-26982-7_7

Miiggel, H., Rhol, T., Speicherl, D., Bihlerl, P., & Cremersl, A. B. (2007). Programming for context-based adaptability: Lessons learned about OOP, SOA, and AOP. *Communication in Distributed Systems*, *15*, 1–12.

Miller, H. J. (2004). Activities in space and time. In D. A. Hensher, K. J. Button, K. E. Haynes & P. R. Stopher, (Eds.), *Handbook of transport geography and spatial systems 5* (pp. 647–660). Emerald Group Publishing. https://doi.org/10.1108/9781615832538-036

Miller, H., Wu, Y.-H., & Miller, H. J. (2001). Computational tools for measuring space-time accessibility within dynamic flow transportation networks; computational tools for measuring space-time accessibility within transportation networks with dynamic flow. *Journal of Transportation and Statistics*, *4*(2). https://www.researchgate.net/publication/2486027

Naeem, M. A., Yahui, M., & Abrar, A. (2025). Optimizing data retrieval latency in IoT through information centric in-network caching. *Alexandria Engineering Journal*, *111*, 468–481. https://doi.org/10.1016/J.AEJ.2024.10.067

Newman, G., Zimmerman, D., Crall, A., Laituri, M., Graham, J., & Stapel, L. (2010). User-friendly web mapping: Lessons from a citizen science website. *International Journal of Geographical Information Science*, *24*(12), 1851–1869. https://doi.org/10.1080/13658816.2010.490532

Norman, D. (2013). *The design of everyday things.* Basic Books.

Oliveri, M., & Reichenbacher, T. (2021). Suitability of colour hue, value, and transparency for geographic relevance encoding in mobile maps. *Abstracts of the International Cartographic Association*, *3*, 14–18. https://doi.org/10.5194/ica-abs-3-224-2021

Onu, F. U., Ikedilo, O. E., Nwoke, B. O., & Okafor, P. (2015). The importance of object-oriented programming in this era of mobile application development. *IOSR Journal of Computer Engineering*, *17*(6), 30–40. https://doi.org/10.9790/0661-17613040

Packia, A. P. A. A., Babu, M., Al-Khalidi, A., Abdulhasan, M. M., Erudiyanathan, D., & Naqvi, A. (2024). (NBCI) Natural brain (human) and computer interaction: Development of user experience with computers through adaptive interface method. *2024 4th International Conference on Advance Computing and Innovative Technologies in Engineering, ICACITE 2024*, 1135–1139. https://doi.org/10.1109/ICACITE60783.2024.10617342

Papadimitriou, E., Theofilatos, A., Yannis, G., Sardi, G. M., & Freeman, R. P. J. (2012). Road safety attitudes and perceptions of pedestrians in Europe. *Procedia - Social and Behavioral Sciences*, *48*, 2490–2500. https://doi.org/10.1016/J.SBSPRO.2012.06.1220

Raper, J. (2007). Geographic relevance. *Journal of Documentation*, *63*(6), 836–852. https://doi.org/10.1108/00220410710836385

Rashid, A. M., Albert, I., Cosley, D., Lam, S. K., McNee, S. M., Konstan, J. A., & Riedl, J. (2002). Getting to know you. *Proceedings of the 7th International Conference on Intelligent User Interfaces*, 127–134. https://doi.org/10.1145/502716.502737

Reichenbacher, T. (2004). *Mobile cartography-adaptive visualisation of geographic information on mobile devices* [Dissertation, University of Zurich]. https://www.researchgate.net/publication/200622119_Mobile_Cartography_-_Adaptive_Visualisation_of_Geographic_Information_on_Mobile_Devices

Reichenbacher, T. (2005). Adaptive egocentric maps for mobile users. *Map-Based Mobile Services: Theories, Methods and Implementations*, 141–158. https://doi.org/10.1007/3-540-26982-7_10

Reichenbacher, T. (2007). The concept of relevance in mobile maps. *Location Based Services and TeleCartography*, 231–246. https://doi.org/10.1007/978-3-540-36728-4_18

Reichenbacher, T. (2019). *Re-visiting fundamental principles of mobile cartography* [Paper presentation]. SIGCHI Conference. https://use.icaci.org/wp-content/uploads/2018/11/Reichenbacher.pdf

Reichenbacher, T., & Bartling, M. (2023). Adaptivity as a key feature of mobile maps in the digital era. *Frontiers in Communication*, *8*. https://doi.org/10.3389/fcomm.2023.1258851

Reichenbacher, T., & De Sabbata, S. (2011). Geographic relevance-different notions of geographies and relevancies. *SIGSPATIAL Special Issues on Geographic Information Retrieval*, *3*, 67–70. https://doi.org/10.5167/uzh-51363

Reichenbacher, T., De Sabbata, S., Purves, R. S., & Fabrikant, S. I. (2016). Assessing geographic relevance for mobile search: A computational model and its validation via crowdsourcing. *Journal of the Association for Information Science and Technology*, *67*(11), 2620–2634. https://doi.org/10.1002/asi.23625

Rock, I., & Palmer, S. (1990). The legacy of Gestalt psychology. *Scientific American*, *263*(6), 84–91. https://doi.org/10.2307/24997014

Roth, R. E. (2017). Visual variables. In *International encyclopedia of geography* (pp. 1–11). Wiley. https://doi.org/10.1002/9781118786352.wbieg0761

Roth, R. E. (2019, July 12). What is mobile first cartographic design? *ICA Workshop on Mobile Map User Experience Design*.

Sarjakoski, L. T., & Nivala, A. M. (2005). Adaptation to context: A way to improve the usability of mobile maps. In *Map-based Mobile Services: Theories, Methods and Implementations* (pp. 107–123). Berlin/Heidelberg: Springer. https://doi.org/10.1007/3-540-26982-7_8

Savino, G. L., Sturdee, M., Rundé, S., Lohmeier, C., Hecht, B., Prandi, C., Nunes, N. J., & Schöning, J. (2021). MapRecorder: Analysing real-world usage of mobile map applications. *Behaviour and Information Technology*, *40*(7), 646–662. https://doi.org/10.1080/0144929X.2020.1714733

Seow, S. C. (2008). Designing and engineering time: The psychology of time perception in software. Indianapolis, IN: Addison-Wesley Professional.

Shu, D., Peng, Y., Zhang, Z., Shi, R., Wu, C., Gan, D., & Li, X. (2024). Distance decay of urban park visitation: Roles of personal characteristics and visitation patterns. *Forests*, *15*(9). https://doi.org/10.3390/F15091589

Steichen, B., & Fu, B. (2019). Towards adaptive information visualization: A study of information visualization aids and the role of user cognitive style. *Frontiers in Artificial Intelligence*, *2*(22), 1–10. https://doi.org/10.3389/frai.2019.00022

Sun, X., Chen, X., Liu, Y., Grundy, J., & Li, L. (2023). LazyCow: A lightweight crowdsourced testing tool for taming Android fragmentation. Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering. https://doi.org/10.1145/3611643.3613098

Swienty, O., Reichenbacher, T., Reppermund, S., & Zihl, J. (2008). The role of relevance and cognition in attention-guiding geovisualisation. *The Cartographic Journal*, *45*(3), 227–238. https://doi.org/10.1179/000870408X311422

Teixeira, G. P., Da Cunha, N. B., Azeredo, C. M., Rinaldi, A. E. M., & Crispim, C. A. (2024). Eating time variation from weekdays to weekends and its association with dietary intake and BMI in different chronotypes: findings from National Health and Nutrition Examination Survey (NHANES) 2017–2018. *British Journal of Nutrition*, *131*(7), 1281–1288. https://doi.org/10.1017/S0007114523002738

Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography*, *46*, 234. https://doi.org/10.2307/143141

Tufte, E. R., & Graves-Morris, P. R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press, *2*(9), 53–121.

Ul Haque, E., Khan, M. M. H., & Fahim, M. A. Al. (2023). The nuanced nature of trust and privacy control adoption in the context of Google. *Conference on Human Factors in Computing Systems - Proceedings*. https://doi.org/10.1145/3544548.3581387/SUPPL_FILE/3544548.3581387-VIDEO-PREVIEW.MP4

van Elzakker, C. P. J. M., & Ooms, K. (2017). Understanding map uses and users. In *The Routledge handbook of mapping and cartography* (pp. 55–67). Routledge. https://doi.org/10.4324/9781315736822-6

Wilson, P. (1973). Situational relevance. *Information Storage and Retrieval*, *9*(8), 457–471. https://doi.org/10.1016/0020-0271(73)90096-X

Wingrave, C. A., Bowman, D. A., & Ramakrishnan, N. (2001). *Affordances and feedback in nuance-oriented interfaces*. Department of Computer Science, Virginia Polytechnic Institute & State University http://hdl.handle.net/10919/20048

Xhafa, F. (2018). *Lecture notes on data engineering and communications technologies series editor*. https://doi.org/https://doi.org/10.1007/978-3-319-67925-9

Xu, Z., Chen, Z. D., & Nie, H. (2014). Handheld computers: Smartphone-centric wireless applications. *IEEE Microwave Magazine*, *15*(2), 36–44. https://doi.org/10.1109/MMM.2013.2296213

Yin, H., Cui, B., Huang, Z., Wang, W., Wu, X., & Zhou, X. (2015). Joint modeling of users' interests and mobility patterns for point-of-interest recommendation. *MM 2015 - Proceedings of the 2015 ACM Multimedia Conference*, 819–822. https://doi.org/10.1145/2733373.2806339

Yu, D., Xu, K., Wang, D., Yu, T., & Li, W. (2019). Point-of-interest recommendation based on user contextual behavior semantics. *International Journal of Software Engineering and Knowledge Engineering*, *29*(11–12), 1781–1799. https://doi.org/10.1142/S0218194019400217

Yu, W., & Xie, F. (2024). Research on traffic congestion prediction based on XGBoost. *Frontiers in Traffic and Transportation Engineering*, *4*(1). https://doi.org/10.23977/FTTE.2024.040101

Zhang, D., Adipat, B., & Mowafi, Y. (2009). User-centered context-aware mobile applications-the next generation of personal mobile computing. *Communications of the Association for Information Systems*, *24*(1), 27–46. https://doi.org/10.17705/1cais.02403

Zhu, L., Shu, S., & Zou, L. (2022). XGBoost-based travel time prediction between bus stations and analysis of influencing factors. *Wireless Communications and Mobile Computing*, *2022*. https://doi.org/10.1155/2022/3504704

Zingaro, D., & Reichenbacher, T. (2022a). Exploratory analysis of mobile app usage in relation to distance from home. *International Conference on Location Based Services*, 59–67. https://doi.org/10.5167/uzh-224336

Zingaro, D., & Reichenbacher, T. (2022b). Modelling and communicating geographic relevance in a digitally transformed world using a digital twin. *Abstracts of the ICA*, *5*, 1–2. https://doi.org/10.5194/ica-abs-5-123-2022

Zingaro, D., Reichenbacher, T., Bartling, M., & Fabrikant, S. I. (2024). Exploring the relation between sense of direction and spatial anxiety in everyday mobile map app use. *Leibniz International Proceedings in Informatics, LIPIcs*, *315*. https://doi.org/10.4230/LIPICS.COSIT.2024.15

Appendices

Appendix 1: Basemap Designs

Day Basemap (Standard)

Outdoor Basemap

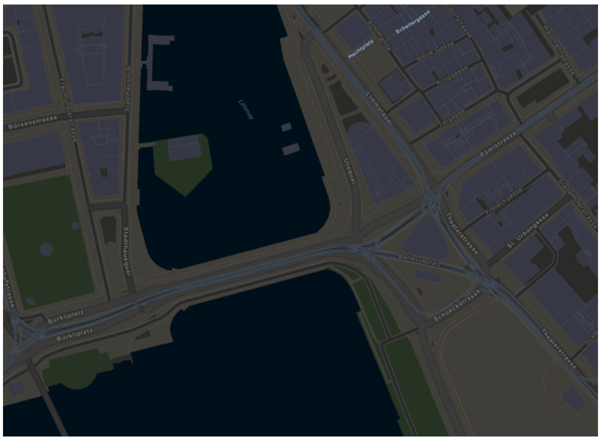

Night Basemap

Food Basemap

Transportation Basemap

Appendix 2: Class Diagram

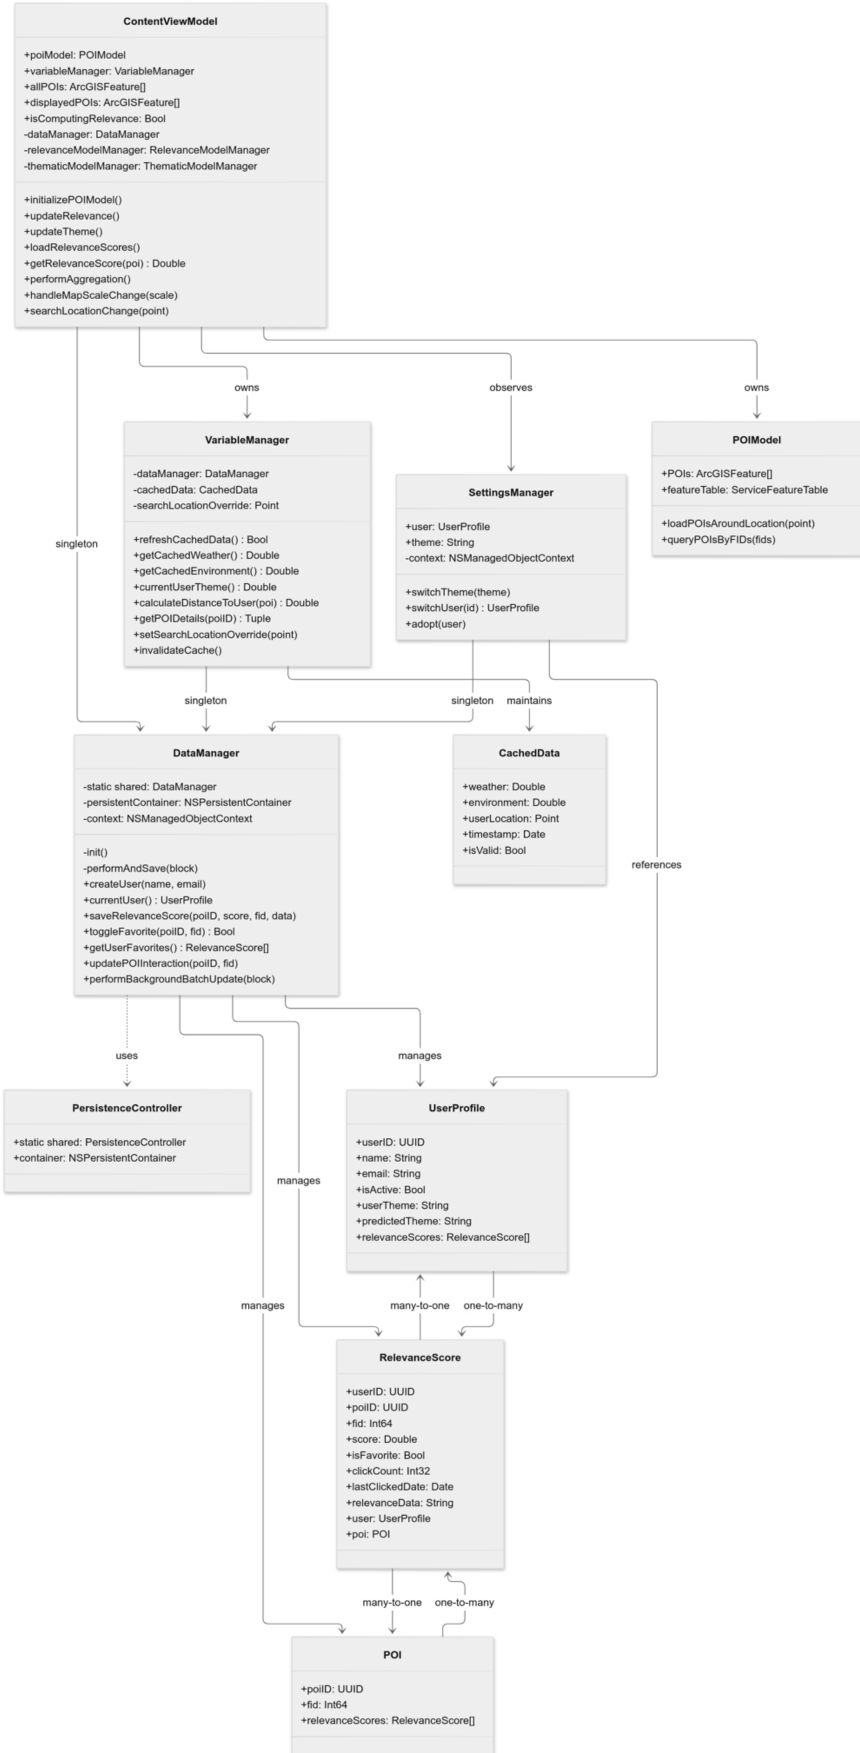

Appendix 3: State Diagram

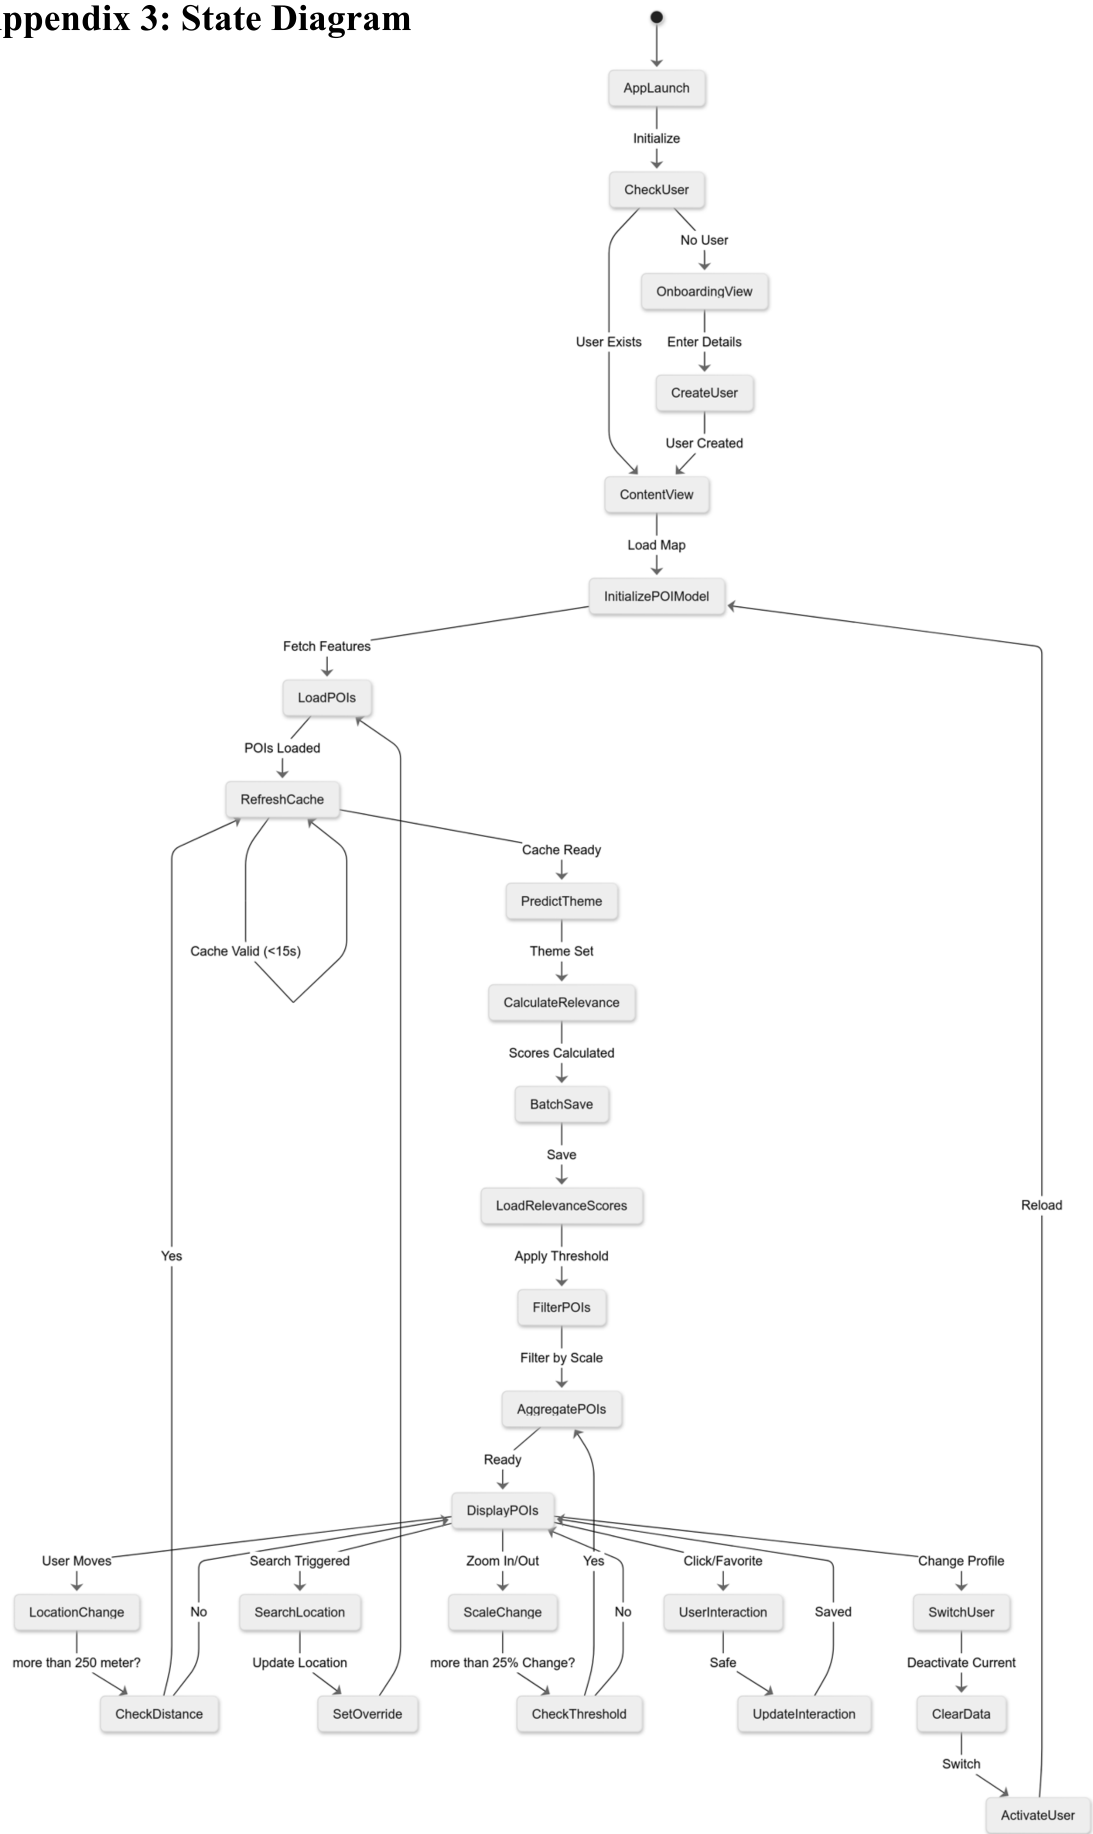

Appendix 4: Amenity Groupings of Point of Interests

Culture

- artwork, arts_centre
- bar, pub
- museum
- theatre
- nightclub
- cinema
- biergarten
- theme_park
- community_centre, public_building, town_hall

Health

- doctors, pharmacy
- veterinary
- dentist
- hospital, clinic
- nursing_home

Education

- school, college, university
- kindergarten
- library

General

- post
- car_sharing, general, comms_tower, hunting_stand, water_works, observation_tower, wastewater_plant, water_tower, embassy, prison, courthouse, car_wash, graveyard
- shelter, chalet
- bank
- police, fire_station
- toilet
- atm
- vending_any, vending_parking, vending_machine, vending_cigarette
- telephone
- recycling_metal, recycling_paper, recycling, recycling_glass, recycling_clothes
- waste_basket
- drinking_water, fountain, water_well

Transportation

- airport, airfield
- bus_station, bus_stop
- helipad
- taxi
- ferry_terminal
- railway_halt, railway_station
- tram_stop

Nature

- picnic_site
- park, dog_park
- zoo
- peak
- beach
- spring

Sports

- golf_course
- sports_centre, playground, pitch, ice_rink, track, stadium
- swimming_pool

Tourism

- attraction, fort, battlefield, windmill, lighthouse, water_mill, monument, ruins, castle, archaeological, memorial, wayside_cross, cliff
- viewpoint, tower
- tourist_info, wayside_shrine

Religious

- jewish, buddhist, christian_lutheran, hindu, christian_anglican, muslim_sunni, christian_methodist, muslim, christian_protestant, christian_evangelical, christian_catholic, christian

Shops

- supermarket, convenience, beverages, market_place, mall, greengrocer
- hairdresser
- clothes
- bakery, beauty_shop, car_dealership, florist, sports_shop, jeweller, butcher, shoe_shop, furniture_shop, gift_shop, doityourself, mobile_phone_shop, car_rental, toy_shop, computer_shop, department_store, garden_centre, outdoor_shop, video_shop, travel_agent, laundry, stationery, chemist
- bicycle_shop, bicycle_rental
- kiosk, newsagent
- bookshop
- optician

Restaurants

- restaurant, fast_food, food_court
- cafe

Hotel

- camp_site, caravan_site
- hotel, guesthouse, alpine_hut, hostel, motel

Appendix 5: Core Data Sample

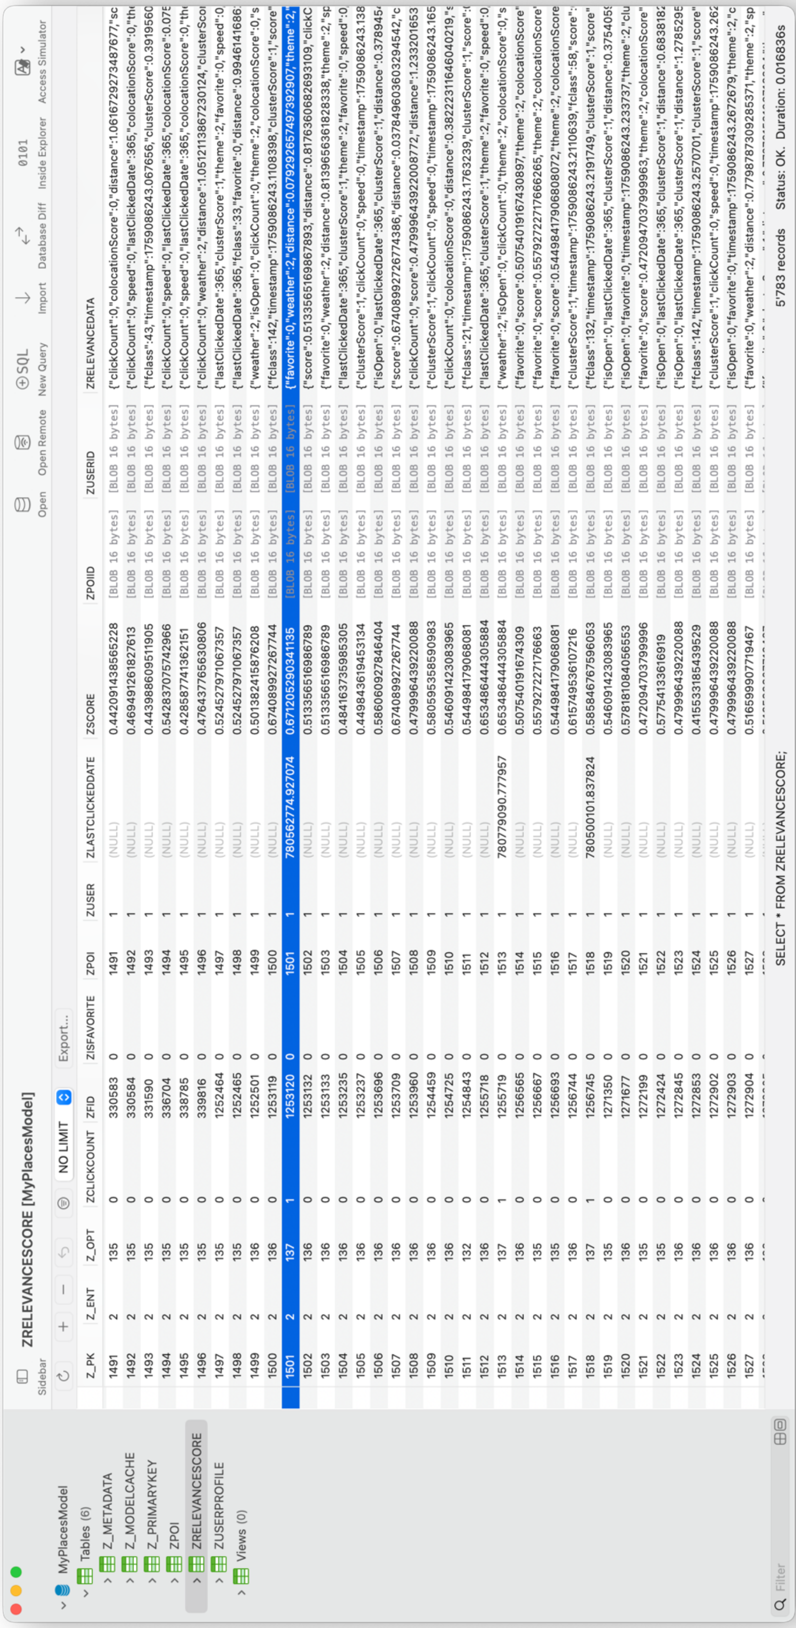

Appendix 6: Computational Application Results

timestamp	environment	poi_count	displayed_pois	reduction_rate	scale	relevance_time_ms	thematic_time_ms	aggregation_time_ms	loading_time_ms	total_time_ms	relevance_per_poi_ms	total_time_s
2025-09-29T15:16:12Z	urban	1587	425	0.73	500	1481.94	1103.56	5110.26	1405.55	9101.31	0.934	9.10131
2025-09-29T15:16:21Z	urban	1587	318	0.8	1500	1652.11	1.15	5365.36	661.6	7680.22	1.041	7.68022
2025-09-29T15:16:30Z	urban	1587	66	0.96	5000	1630.49	649.41	5806.22	626.65	8712.77	1.027	8.71277
2025-09-29T15:16:32Z	rural	116	69	0.41	500	257.84	1.08	270.41	778.11	1307.44	2.223	1.30744
2025-09-29T15:16:33Z	rural	116	60	0.48	1500	280.74	2.21	307.49	208.15	798.6	2.42	0.7986
2025-09-29T15:16:34Z	rural	116	27	0.77	5000	267.64	2.32	412.21	198.45	880.63	2.307	0.88063
2025-09-29T15:16:35Z	outdoor	9	6	0.33	500	61.16	2.55	20.99	321.12	405.81	6.795	0.40581
2025-09-29T15:16:36Z	outdoor	9	6	0.33	1500	62.41	2.73	20.63	178.08	263.84	6.934	0.26384
2025-09-29T15:16:37Z	outdoor	9	3	0.67	5000	53.66	2.52	32.11	584.12	672.41	5.962	0.67241

Declaration of Independence

Personal declaration: I hereby declare that the submitted thesis is the result of my own, independent work. All external sources are explicitly acknowledged in the thesis.
Additionally, Grammarly was used to check the grammar and enhance the clarity of the text. The AI models Claude (Anthropic) and ChatGPT (OpenAI) were used as debugging assistants to help find errors and performance issues in the implemented application code.

Johannes P. Guler, 30.09.2025