TU WIEN Informatics

Visualization of Points of Interest in 3D Digital Maps

BACHELORARBEIT

zur Erlangung des akademischen Grades

Bachelor of Science

im Rahmen des Studiums

Medieninformatik und Visual Computing

eingereicht von

Julia Hofmann
Matrikelnummer 11827004

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Associate Prof. Dr.in techn. Manuela Waldner, MSc
Mitwirkung: Univ.Ass.in Angeliki Grammatikaki, MSc

Wien, 22. September 2025

Julia Hofmann   Manuela Waldner

TU WIEN Informatics

Visualization of Points of Interest in 3D Digital Maps

BACHELOR'S THESIS

submitted in partial fulfillment of the requirements for the degree of

Bachelor of Science

in

Media Informatics and Visual Computing

by

Julia Hofmann
Registration Number 11827004

to the Faculty of Informatics

at the TU Wien

Advisor: Associate Prof. Dr.in techn. Manuela Waldner, MSc
Assistance: Univ.Ass.in Angeliki Grammatikaki, MSc

Vienna, September 22, 2025

Julia Hofmann   Manuela Waldner

Erklärung zur Verfassung der Arbeit

Julia Hofmann

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.
Ich erkläre weiters, dass ich mich generativer KI-Tools lediglich als Hilfsmittel bedient habe und in der vorliegenden Arbeit mein gestalterischer Einfluss überwiegt. Im Anhang „Übersicht verwendeter Hilfsmittel" habe ich alle generativen KI-Tools gelistet, die verwendet wurden, und angegeben, wo und wie sie verwendet wurden. Für Textpassagen, die ohne substantielle Änderungen übernommen wurden, habe ich jeweils die von mir formulierten Eingaben (Prompts) und die verwendete IT-Anwendung mit ihrem Produktnamen und Versionsnummer/Datum angegeben.

Wien, 22. September 2025

Julia Hofmann

Kurzfassung

Orientierungspunkte bzw. Erkennungszeichen sind zentral für die räumliche Orientierung, doch ihre Gestaltung in ländlichen 3D-Geländekarten ist nicht ausreichend erforscht. Diese Arbeit untersucht, wie Orientierungspunkte visualisiert werden sollten, um Erforschung und Verständnis mit Fokus auf ländliche 3D-Karten zu unterstützen. Ein modularer, webbasierter Prototyp wird vorgestellt, der solche “Points of Interest (PoIs)” aus OpenStreetMap mittels einer asynchronen Vorverarbeitungspipeline bereinigt, kategorisiert und für den Client partitioniert. Eine quadtree-gestützte Datenverwaltung ermöglicht eine dynamische Level-of-Detail (LoD) Strategie, welche die Dichte der Orientierungspunkte reguliert, um die Lesbarkeit über verschiedene Distanzen zu erhöhen. Die verschiedenen Visualisierungsmodi (Textbeschriftungen, abstrakte Symbole, piktografische Icons, 3D-Modelle und Heatmaps) werden in eine 3D-Geländekarte implementiert und können zur Laufzeit verglichen werden. Nutzerzentrierte Interaktionsmöglichkeiten wie Kategoriefilter ergänzen die Darstellung.

Die Arbeit verbindet Systemgestaltung mit einer literaturbasierten Analyse, um die jeweiligen Abwägungen zwischen den Modi präzise herauszuarbeiten. Die Ergebnisse zeigen maßstabsabhängige Vor- und Nachteile der Visualisierungsmodi, heben die Kontrolle der Objektdichte vor der Einführung zusätzlicher Details hervor sowie eine auf Aufgabe und Maßstab abgestimmte Balance zwischen Abstraktion und Realismus. Die Beiträge sind ein funktionsfähiger technischer Rahmen, der Orientierungspunktvisualisierung für ländliche 3D-Gelände operationalisiert, und eine strukturierte Synthese, die verdeutlicht, wann und warum bestimmte Darstellungen vorteilhaft sind. Die Evaluation erfolgt theoretisch und skizziert Implikationen und Hypothesen für künftige empirische Validierung.

Abstract

Landmarks are central to spatial orientation, yet the design of landmarks in rural environments has not been investigated so far. This thesis investigates how landmark visualization modes should be designed to support spatial exploration and landmark comprehension on rural 3D maps. A modular, web-based prototype integrates Points of Interest (PoIs) from OpenStreetMap in an asynchronous preprocessing pipeline that cleans, categorizes, and partitions data for client-side use. A quadtree-based data management approach provides dynamic level-of-detail (LoD) to regulate density and maintain legibility across scales. Within this framework, interchangeable visualization modes (text labels, abstract symbols, pictorial icons, 3D models, and heatmaps) are implemented on a 3D terrain map of Austria, complemented by category-based filtering and details-on-demand interactions.

The thesis combines system design with a literature-grounded analysis to articulate the trade-offs among these modes. The resulting guidance emphasizes scale-dependent staging of encodings, control of density before the introduction of detail, and a task- and scale-sensitive balance between abstraction and realism, with semantics exposed through lightweight interaction rather than persistent annotation. The contributions are twofold: a functional technical framework that operationalizes landmark visualization for rural 3D terrain, and a structured synthesis that clarifies when and why particular visualization methods are advantageous. Evaluation proceeds theoretically rather than through user studies, and the thesis outlines implications and hypotheses for future empirical validation.

Contents

CHAPTER 1

Introduction

In rural outdoor settings, accurate and user-centered landmark visualization on three-dimensional (3D) maps is essential for efficient navigation and spatial orientation. Landmarks function as salient reference points that anchor users' mental representations of space and inform wayfinding decisions, particularly where other cues, such as signage or dense infrastructure, are sparse [Yesiltepe et al., 2021, Deakin, 1996]. Prior technical work has demonstrated the feasibility of automatically extracting and evaluating landmarks from cadastral and surface data, enabling their systematic integration into digital navigation systems [Elias and Brenner, 2005].

Although 3D terrain maps provide accurate depictions of topography, they often lack strategies for encoding landmarks in cognitively supportive ways. Research shows that higher visual fidelity does not necessarily improve navigation performance. In fact, abstract or stylized depictions can often facilitate route learning and reduce cognitive load compared to photorealistic renderings [Çöltekin et al., 2018, Smallman and Cook, 2011]. At the same time, recent controlled experiments show that realistic 3D landmarks can enhance visual attention and spatial learning in mobile navigation, particularly for users with lower spatial ability [Kapaj et al., 2022, Kapaj et al., 2024]. This duality highlights the need to balance visual realism and abstraction in landmark visualization carefully. Consistent with prior work, landmark depictions range along a spectrum between realistic and abstract forms [Elias and Paelke, 2008].

Building on this premise, the thesis investigates landmark representation tailored to rural terrain, integrating dynamic Level of Detail (LoD) strategies and evaluating design choices through a human-centered lens. LoD denotes scale-dependent generalization/selection that controls how many and which features are shown at a given viewing scale. The subsequent chapters detail the conceptual background, a prototype that operationalizes category-based 2D/3D landmark visualizations on 3D terrain, and a discussion of empirical findings and limitations that inform recommendations for future work.

1.1 Motivation and Context

Landmarks are salient Points of Interest (PoIs) that help individuals interpret and navigate their surroundings [Deakin, 1996]. A Point of Interest (PoI) is any geographically referenced object that is meaningful for navigation or analysis. Their relevance is pronounced in rural regions, where street signage and other infrastructural cues are sparse. Integrating landmark data from crowdsourced geospatial platforms offers a practical response to this challenge. OpenStreetMap (OSM), the most prominent example of volunteered geographic information, provides a rich, global dataset of natural and cultural features [OpenStreetMap Foundation, 2025]. While its accessibility and coverage are advantageous, positional and thematic accuracy vary across regions [Al-Bakri, 2015]. For applications where absolute precision is not critical, OSM still provides a robust foundation for landmark-focused visualization and navigation research.

Prior research examines how cartographic symbol choices shape attention and memory for landmarks. An eye-tracking study comparing text, icon, and vignette symbols reported similar fixation metrics across classes, with text yielding the highest average recall [Franke and Schweikart, 2017]. Complementary findings indicate that overall visual complexity strongly influences attention allocation and route memory, with higher complexity often reducing fixation on task-relevant features [Keil et al., 2020]. Cognitive load research further suggests an optimal density of landmark information: too few landmarks hinder learning, whereas too many can overload limited working memory [Cheng et al., 2022]. In addition, studies in virtual navigation show that higher visual realism does not necessarily improve performance, and that individual differences in spatial ability are consequential [Çöltekin et al., 2018, Smallman and Cook, 2011]. These insights argue for careful, evidence-based design rather than assuming that greater realism is inherently beneficial.

Foundational system-oriented work demonstrates the feasibility of automatically extracting candidate landmarks from cadastral and surface data and assessing their visibility with digital surface models [Elias and Brenner, 2005]. Building on this technical foundation and the human factors evidence, the present thesis focuses on rural terrain settings and employs an asynchronous preprocessing pipeline to curate, categorize, and filter OSM-derived landmarks. The resulting dataset supports category-based two-dimensional (2D)/3D encodings on 3D terrain with LoD controls aimed at maintaining legibility and managing cognitive load.

1.2 Problem Statement

Despite these advances, it remains unclear which landmark visualization techniques are most effective for supporting navigation in rural outdoor contexts. The existing studies largely stem from abstract 2D maps, virtual environments, or urban settings, leaving rural 3D environments underexplored [Yesiltepe et al., 2021]. In particular, findings on symbol classes, visual complexity, or realism do not directly transfer to interactive outdoor terrain maps that combine natural and cultural landmarks.

This problem is especially pronounced in rural settings. Unlike urban areas, where navigation can rely on dense networks of streets and signage, rural navigation often depends more strongly on landmarks for spatial orientation [Elias and Brenner, 2005]. At the same time, systematic reviews point out the lack of design guidance on how such landmarks should be visualized in different contexts [Yesiltepe et al., 2021]. The trade-offs between different visualization methods, from abstract symbols to detailed 3D models, remain poorly understood in outdoor rural terrain.

To address this gap, this thesis implements a flexible prototype for rural 3D maps and provides a structured analysis of landmark visualization modes, assessing their respective strengths and limitations against prior research.

1.3 Objectives and Scope

The thesis examines how distinct landmark visualization modes support spatial exploration and landmark comprehension in rural 3D map environments. To address this, the prototype integrates PoIs derived from OSM into a web-based 3D terrain map of Austria. An asynchronous preprocessing pipeline curates and classifies OSM data for use in multiple visualization modes (text labels, abstract symbols, pictorial icons, heatmaps, and 3D models) to enable comparison of their characteristics within a single system. The work combines technical implementation with a literature-grounded analysis to position these modes within existing research.

More specifically, the thesis pursues the following objectives:

- Design a preprocessing pipeline that transforms raw OSM extracts into a structured, categorized dataset suitable for web-based 3D visualization.
- Implement a quadtree-based data management approach that supports scalable rendering and a dynamic Level of Detail strategy for large numbers of landmarks.
- Realize multiple visualization layers (text labels, symbols, pictorial icons, 3D models, and heatmaps) to enable direct, within-system comparison.
- Provide User Interface (UI) controls for category-based filtering, interactive exploration, and access to landmark details in a 3D context.
- Analyze the relative strengths and limitations of the implemented visualization modes through systematic comparison with findings from prior literature.

The scope is deliberately constrained. The prototype is limited to the national extent of Austria to maintain manageable data volumes while covering diverse rural contexts. The focus is on landmark visualization rather than route planning, guidance, or navigation algorithms. Real-time data ingestion was not possible because of performance concerns with classification, so landmark data is prepared via asynchronous preprocessing.

Evaluation is conducted as a theoretical comparison with prior studies rather than an empirical user study. Within these boundaries, this thesis contributes both a technical framework and a structured analysis of landmark visualization strategies tailored to rural 3D environments.

1.4 Main Research Question

The previous sections outlined the challenges of landmark visualization in rural navigation and established the scope and objectives of this thesis. While prior research has examined symbol classes, visual complexity, and realism, their findings do not fully transfer to rural 3D map environments. This work therefore consolidates the identified issues into a single guiding perspective. Accordingly, the central research question is:

> How should landmark visualization modes be designed to support spatial exploration and landmark comprehension in rural 3D terrain maps?

1.5 Contributions of This Thesis

The contributions are technical and conceptual, combining a functional prototype with a structured discussion of landmark visualization strategies for rural 3D contexts. Technically, the work delivers a modular, web-based system that integrates OSM-derived data through an asynchronous preprocessing pipeline and exposes multiple visualization modes within a single interactive environment. Conceptually, the thesis positions these modes against existing research to articulate design trade-offs relevant to rural terrain.

Data Preprocessing Pipeline

A structured pipeline cleans, classifies, and spatially organizes OSM data for landmark visualization. The outputs are partitioned and indexed to support scalable client-side use, ensuring that the varying volumes of landmarks present in rural settings can be handled without loss of responsiveness.

Data Management and LoD

A quadtree-based approach underpins scalable delivery and rendering. The dynamic LoD strategy selectively exposes landmarks by category and scale, reducing visual clutter and helping maintain legibility during navigation and exploration on 3D terrain.

Visualization Layers and Interaction

The prototype introduces six alternative visualization layers that can be switched dynamically during runtime: text labels, simplified symbols, pictorial icons, 3D models, combined 2D and heatmaps. Complementary UI elements provide category-based filtering and information popups, supporting targeted exploration and inspection of landmark attributes in a 3D context.

Analytical Framing

An analytical framework synthesizes the relative strengths and limitations of each visualization mode with respect to findings in prior literature. The evaluation proceeds as a structured, theory-informed comparison rather than an empirical user study, emphasizing design implications for rural 3D environments.

Together, these contributions provide a concrete technical foundation for landmark visualization on 3D terrain and a focused analysis that clarifies when and why particular modes are advantageous for spatial exploration and landmark comprehension in rural settings.

1.6 Thesis Structure

The remainder of this thesis moves from context to contribution. Chapter 2 establishes the theoretical and technical background in 3D cartographic visualization, landmark representation, and OSM-based spatial data integration, and distills the research gaps that motivate this work. Chapter 3 presents the concept, and Chapter 4 operationalizes these ideas. Chapter 5 interprets the results with respect to the research question and situates them within related work, while also reflecting on limitations and outlining key design implications, and finally concludes by summarizing the contributions and outlining directions for future work.

CHAPTER 2

Background and Related Work

This chapter provides an overview of the relevant research and theoretical foundation that support this thesis. It focuses on how landmark representation (Section 2.2) and cognitive principles (Section 2.3) influence the design and usability of 3D maps, with particular emphasis on rural and mountainous navigation contexts where landmark availability and visibility differ fundamentally from urban environments. While existing research has largely focused on urban contexts with dense, architectural landmarks, rural and mountainous environments introduce additional challenges such as sparse point distributions, natural landmark reliance, and elevation-based occlusion.

Modern 2D mapping applications frequently rely on conventional landmark representations such as standardized icons or text labels. These approaches, originally designed for flat, 2D contexts, often fail to address the perceptual demands posed by 3D environments. [Franke and Schweikart, 2017] demonstrated that landmark visualization (text labels, icons, vignettes) systematically shapes users' mental representations on 2D maps, with text yielding the highest average recall. However, these findings do not directly address 3D conditions such as occlusion or depth ordering [Yesiltepe et al., 2021].

In 3D contexts, visibility is a prerequisite for salience and use, and managing added information is critical. However, the effect is not monotonic: studies show that including landmarks can aid learning without raising cognitive load, while excessive numbers may increase it [Yesiltepe et al., 2021, Cheng et al., 2022]. Beyond symbol choice, 3D landmark placement must handle occlusion, viewpoint changes, and depth ordering, motivating dynamic selection and placement strategies rather than static 2D conventions. [Cliburn et al., 2007]

Similarly, Keil and Edler show that visual complexity modulates attentional allocation: under higher complexity, participants concentrated attention closer to routes and decision points. Importantly, this did not produce clear detriments to route memory in their recognition task [Keil et al., 2020]. While realistic landmark depictions are often perceived as more attractive, they do not reliably improve task performance. In mobile navigation, realistic 3D landmarks drew more attention but did not yield faster or more accurate performance [Kapaj et al., 2024]. Laboratory studies likewise show that highly realistic map displays can impose greater cognitive demands than abstract representations [Çöltekin et al., 2018].

In parallel with perceptual concerns, there is also a methodological gap in how landmarks are selected and integrated into 3D systems. Carbonell Carrera et al. emphasize the pedagogical benefits of 3D terrain representations for map-reading skill development, reinforcing the value of spatially rich displays in education [Carbonell Carrera et al., 2017]. Yet systematic comparisons of landmark visualization techniques beyond predominantly urban, 2D settings remain limited [Yesiltepe et al., 2021]. In contrast, Elias and Brenner present a road-junction–oriented, DSM-based visibility pipeline for automatic selection of building landmarks [Elias and Brenner, 2005]. Subsequent reviews note important limitations for broader contexts, for example, reliance on building features and largely single-view visibility, which complicate transfer to rural, natural-feature landscapes [Yesiltepe et al., 2021]. This raises the open question of how well such techniques generalize to rural landscapes, where natural features and sparsely distributed structures constitute the primary navigational cues.

2.1 Web-Based 3D Terrain Visualization

3D terrain modeling plays a critical role in enhancing spatial comprehension and supporting navigation, especially in rural and outdoor environments, where elevation and landform features strongly influence visibility and movement constraints. High-fidelity 3D terrain representations have been shown to support spatial comprehension and learning tasks, underlining their value in both educational and navigational settings [Carbonell Carrera et al., 2017].

Digital Elevation Models (DEMs). 3D terrain visualization combines elevation data and imagery into an interactive surface that supports navigation and spatial reasoning. DEMs provide the structural basis for simulating terrain surfaces in geospatial applications. Acquired through remote sensing techniques such as LiDAR or photogrammetry, DEMs encode terrain elevation into a gridded matrix of height values (Figure 2.1a). The resolution and vertical accuracy of a DEM directly affect both the realism of rendered terrain and the clarity of overlaid landmark features. Elias and Brenner, for example, leveraged detailed digital surface models to analyze landmark visibility [Elias and Brenner, 2005], demonstrating how terrain modeling supports effective route planning and spatial decision-making.

Accurate terrain representation is central to landmark-based navigation, since visibility and contextual salience depend on both position and surrounding topography. Figure 2.1 outlines the construction workflow: orthophoto tiles (Figure 2.1b) provide aerial imagery produce a terrain surface closely resembling the real environment when draped over the

(a) DEM [Maps-for-Free, 2025]. (b) Orthophoto Tiles. (c) Finished terrain.

Figure 2.1: Constructing the terrain surface in a web-based 3D map.

DEM (Figure 2.1c). Terrain modelling and landmark visualization are therefore designed together to maintain visual coherence and spatial legibility across scales.

Browser-Based Visualization Frameworks Several browser-oriented frameworks support web-based terrain. CesiumJS offers globe-scale terrain streaming and 3D Tiles but introduces a relatively heavy abstraction and runtime footprint [Cesium GS, Inc., 2025]. Mapbox GL JS emphasizes vector tiles and integrates terrain and imagery in a map-centric workflow suited to PoI overlays [Mapbox, 2025]. Deck.gl provides a flexible WebGL layer model for large datasets, though advanced use benefits from graphics expertise [OpenJS Foundation, 2025]. MapLibre GL JS, the community successor to Mapbox GL JS v1, renders vector tiles in the browser, supports terrain via raster DEMs, and provides an open plugin ecosystem without vendor lock-in [MapLibre, 2025].

Based on these considerations, Mapbox GL JS was selected [Mapbox, 2025]. The decisive factor is its map-centric *source/layer* model, which cleanly separates data ingestion from visualization and thus supports multi-representation PoI overlays from shared sources using data-driven filters and expressions. This framework supports both 2.5D vector extrusion and full 3D terrain draping, which are referred to collectively as '3D,' as it underpins the multi-mode visualization framework, including the 3D model mode described in Section 4.7.

2.2 Landmark Representation in 2D and 3D Maps

Landmarks are fundamental to navigation, serving as anchors for spatial orientation, route memory, and cognitive map construction [Deakin, 1996]. Their importance is well-established in both cognitive and cartographic literature, yet most empirical studies focus on static or 2D representations, predominantly in urban contexts. In contrast, interactive, rural 3D environments introduce new perceptual variables such as occlusion, elevation change, and dynamic viewpoint shifts, which complicate the direct transfer of landmark design principles established in 2D cartography. As 3D maps become more prevalent in web-based applications, the question of how best to represent landmarks in such environments becomes increasingly relevant. Studies emphasize that visual realism, depth perception, and user spatial ability interact in complex ways to affect navigational performance in 3D contexts [Çöltekin et al., 2018, Kapaj et al., 2024].

Iconography and Text Labeling

Landmark visualization methods commonly rely on symbolic icons, textual labels, or photographic imagery. The effectiveness of each approach depends on the task context, user familiarity, and the spatial complexity of the environment. Eye-tracking experiments comparing icons, vignettes, and textual labels in abstract 2D maps suggest that while textual labels require more screen space, they yielded higher recall and supported more efficient mental encoding [Franke and Schweikart, 2017]. Other research found that realistic visual imagery, such as photos of landmarks, aligns better with users' ground-level visual perspectives and enhances landmark recognition in navigation [Lee et al., 2001]. More recent work supports a nuanced view: abstract 3D landmark representations reduce cognitive load in simple environments, while realistic 3D depictions are more effective in complex spatial tasks [Kapaj et al., 2022]. Consequently, the balance between abstraction and realism in landmark visualization must be carefully calibrated in relation to both the map context and the user's specific spatial needs.

Abstraction Levels for Landmark Representation

Depicting landmarks at appropriate Level(s) of Abstraction (LoA) is an established strategy in cartography and 3D geovisualization. Prior work characterises landmark depiction along a continuum from photorealistic imagery through drawings and silhouettes to icons and labels, with the choice governed by landmark function, recognisability, and communicative intent [Elias and Paelke, 2008]. In this view, LoA complements the geometric LoD by controlling thematic granularity and visual stylization so that depiction remains legible under varying viewing conditions.

Within interactive 3D map systems, LoA selection and transitions have been formalised as context-dependent rather than novel mechanisms. Semmo et al. define LoA with respect to viewing distance, task, and display resolution, and demonstrate smooth transitions between abstraction levels by means of animation and blending, coupled with saliency-guided emphasis [Semmo et al., 2012]. Pasewaldt et al. position such transitions within a cartography-oriented rendering pipeline aimed at reducing clutter and improving comprehensibility in complex 3D scenes [Pasewaldt et al., 2012]. These contributions collectively establish LoA selection and transition techniques as standard components for attention management and visual clarity.

The present work adopts these established ideas and operationalises them in an exemplary visualization layer for a rural context. Abstraction levels are mapped to landmark categories following the design space outlined above and are modulated by viewing distance (far, mid, near) so that depiction shifts from more symbolic to more detailed forms as landmarks approach the viewer.

Managing Clutter and Level of Detail (LoD)

Interactive map environments introduce challenges related to visual clutter, which can overwhelm users and reduce navigational clarity. To mitigate this, cartographic systems implement clustering and LoD techniques. Reviews emphasize that reducing visual clutter is essential for wayfinding clarity, yet systematic design guidance for landmark clustering remains limited [Yesiltepe et al., 2021].

LoD strategies adapt the amount and granularity of visible data based on contextual parameters such as zoom level, viewing angle, or user location. Prior work has demonstrated that dynamically placed landmarks at decision points improve spatial memory and reduce navigational errors in virtual environments [Cliburn et al., 2007]. Similarly, an adaptive visualization model that personalizes landmark representations based on the user's spatial familiarity shows significant improvements in navigation accuracy [Zhu et al., 2022]. These findings underscore the value of responsive, context-aware landmark rendering systems that adapt to user intent and environmental complexity.

In summary, landmark representation in interactive 3D environments must address visual, cognitive, and spatial challenges. It requires empirically grounded strategies for managing clutter and implementing adaptive LoD methods. The present thesis responds to this need by evaluating and deploying landmark visualization techniques optimized for rural and mountainous 3D terrains, with an emphasis on adaptability, clarity, and spatial relevance.

2.3 Human-Centered Considerations in Landmark Visualization

While this thesis does not conduct a user study, its design is grounded in established findings on visual attention and landmark perception. Prior research in spatial cognition and digital cartography highlights how specific visualization strategies can support orientation and reduce perceptual overload in map-based tasks.

Managing Visual Complexity through Abstraction and Relevance

Research shows that increasing visual complexity in map interfaces alters attentional allocation, with participants concentrating attention more closely on routes and decision points [Keil et al., 2020]. While this did not impair route memory, it underscores the risks of excessive clutter and supports the need for adaptive filtering strategies, such as scale-dependent landmark visibility. Representation style strongly influences recall, with textual labels yielding higher memory performance, while icons and vignettes offered more immediate perceptual accessibility [Franke and Schweikart, 2017]. These findings suggest that consistent, legible symbolization is critical for supporting efficient mental encoding.

These insights inform the decision to implement an adaptive landmark rendering system in this thesis. Landmark visibility is filtered by zoom level and category to reduce unnecessary complexity and emphasize task-relevant PoIs. This strategy aligns with recommendations to minimize visual load and optimize landmark selection for interpretability, particularly in outdoor and topographically complex environments. However, a comprehensive review found that most empirical studies on landmark perception remain focused on urban and architectural contexts, leaving rural and natural landscapes underexplored [Yesiltepe et al., 2021]. This gap highlights the need to adapt landmark visualization strategies to outdoor environments, where topographic and natural features predominate.

Balancing Realism and Cognitive Salience

While realistic depictions are often attractive and capture attention, they do not consistently improve task performance and can increase cognitive demands [Kapaj et al., 2024, Çöltekin et al., 2018]. This is a form of "naïve realism", where the assumption that more photorealistic visualizations are inherently superior overlooks the cognitive advantages of abstraction [Smallman and Cook, 2011]. Abstract or iconographic representations, by contrast, reduce interpretive effort while still conveying semantic category cues. To balance clarity and context, this project employs simplified, iconographic representations of PoIs, avoiding full photorealism while preserving semantic information such as category and prominence.

This trade-off is particularly relevant in rural terrain, where map elements must remain interpretable even at oblique camera angles and varying altitudes. Oblique, bird's-eye perspectives have been shown to improve recognition and orientation in spatial search tasks [Nurminen and Sirvio, 2021], underscoring the value of clarity and legibility in dynamic outdoor perspectives. By using a restrained visual language with category-coded symbols, the implementation seeks to support perceptual clarity without burdening the user with interpretive effort.

Application to this Thesis

Drawing on these findings, the implemented system emphasizes three principles: maintaining low visual complexity through LoD control [Keil et al., 2020], using consistent symbolic representation to ensure category recognition [Franke and Schweikart, 2017], and favoring abstraction over realism to balance clarity with attentional salience [Kapaj et al., 2024, Çöltekin et al., 2018]. These measures aim to reduce perceptual load and enhance orientation in outdoor 3D terrain, where topographic variability compounds cognitive demands. While previous approaches were largely restricted to urban landmarks such as buildings and junctions [Elias and Brenner, 2005, Yesiltepe et al., 2021], this thesis applies these principles to rural environments, where navigation relies on sparsely distributed natural and infrastructural features.

2.4 Summary of Research Gaps

Despite advances in landmark visualization and web-based 3D mapping, their application in non-urban terrain remains a key limitation. This section identifies these gaps and articulates how the system developed in this thesis addresses them.

Insufficient Adaptation of Level-of-Detail Strategies for Landmarks

Most contemporary 3D mapping platforms incorporate LoD strategies to manage rendering load and avoid visual clutter. However, the application of LoD to semantic map content such as landmarks remains underexplored in cartographic literature. Prior work emphasizes the importance of adaptive visibility management, particularly for landmarks serving navigational functions [Keil et al., 2020, Cliburn et al., 2007]. Yet, these approaches often neglect topographic context and rely on static distance thresholds. Existing strategies are largely tested in urban, building-dominated settings, limiting their transferability to rural landscapes [Yesiltepe et al., 2021].

This thesis implements a dynamic LoD mechanism based on quadtree partitioning, deriving landmark selection from a spatial index rather than ad hoc thresholds. The approach provides deterministic behavior across viewpoints, predictable density control, and straightforward integration with category-aware selection. Comparable techniques in computer graphics and terrain visualization demonstrate the scalability of quadtree-based LoD control [Lindstrom et al., 1996], but the systematic application to semantic cartographic PoIs such as landmarks has received relatively little attention in this context.

Real-Time Integration of Landmarks from OpenStreetMap

The use of OSM as a base for landmark extraction is widespread, but performance bottlenecks remain a critical barrier for real-time applications. Systems relying on Overpass API queries suffer from latency and rate limits, making them unsuitable for interactive 3D environments [Nurminen and Sirvio, 2021]. Furthermore, inconsistencies in OSM's tagging model introduce classification ambiguities [Elias and Brenner, 2005]. These issues are compounded by persistent concerns regarding OSM data quality and consistency across regions, as highlighted in systematic reviews of volunteered geographic information [Al-Bakri, 2015]. Such inconsistencies particularly complicate landmark-based navigation, since perceptually salient features may be omitted or misclassified in community-sourced datasets [Yesiltepe et al., 2021].

To mitigate these issues, this thesis preprocesses OSM extracts asynchronously into quadtree-indexed tiles, avoiding reliance on live Overpass queries and ensuring near-instant runtime access. The application serves these tiles as static assets, so runtime access is near-instant and independent of live APIs. This design prioritizes low-latency, scalable interaction and reproducibility at the cost of immediacy. Freshness is maintained operationally by regenerating snapshots on demand. This approach mirrors vector-tile preprocessing in large-scale web mapping, but is applied here to a landmark-focused layer in a 3D terrain setting.

Category-Based Interaction for 3D Landmark Exploration

Many existing 3D mapping applications support only binary toggles for PoI layers (such as enabling or disabling all amenities), which is insufficient in settings where users seek specific environmental or cultural features. Task-specific landmark selection significantly influences wayfinding performance, but existing implementations remain largely theoretical and confined to 2D contexts [Yesiltepe et al., 2021].

In contrast, this system introduces a refined category-based UI, where landmark types are grouped into user-centric classes (e.g., food, shelter, tourism, nature). These groups can be toggled at runtime, with the UI directly controlling which categories are loaded and how they are rendered. This design supports both exploratory navigation and focused search within sparse rural terrain, offering improved alignment between user goals and visual feedback.

CHAPTER 3

Concept

On a conceptual level, the system turns raw geographic input into view-dependent landmark representations through a structured pipeline. The flowchart (Fig. 3.1) sketches this path and highlights the feedback mechanism: the current view continually determines which data are required and at what LoD they should be rendered. This concept operationalizes the literature gaps identified in Chapter 2. It couples automated, rural-suited PoI extraction from OSM with a scale-aware delivery path and category-oriented interactions. Rather than proposing a new abstraction scheme, it adopts established LoA practices for visibility-constrained terrain settings.

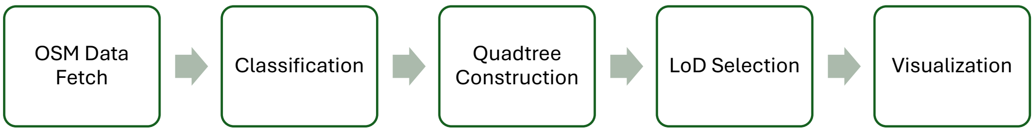

Figure 3.1: Preprocessing and runtime pipeline flowchart.

Starting from the OSM data fetch in Section 3.1, elements for Austria are collected and passed to classification (Section 3.2), where tags are normalized into a consistent category schema used throughout the implementation. The results are written into a spatial index (Section 3.3) that enables efficient retrieval. Given the current camera and scene density, LoD selection decides which representation is appropriate, while delivery provides cacheable chunks. Finally, visualization (Section 3.4) translates these decisions into user-facing representations and lightweight interactions.

3.1 Integration of Geographic Data

Effective landmark visualization depends not only on visual design but also on the accurate integration of underlying geospatial data. OSM is a collaboratively maintained, open-access geospatial database with extensive coverage of spatial features and serves as the source for geographic features required by this system. Each geographic element in OSM is described through flexible key–value tag structures (e.g., `amenity=cafe`, `tourism=attraction`).

Access to OSM data is typically provided via the Overpass API [OpenStreetMap Wiki, 2025]. To ensure reproducibility, predictable performance, and independence from online rate limits, an offline snapshot strategy is adopted. The area of interest is partitioned into a regular grid; results are cached locally to eliminate redundant downloads across runs. The outcome is a reproducible, locally cached snapshot that can be parsed, classified, and indexed without further network calls.

Data Quality Considerations. As a collaboratively maintained dataset, OSM exhibits uneven completeness, heterogeneous tagging, and variable positional accuracy, with pronounced gaps in rural areas [Al-Bakri, 2015]. Names may be missing or multilingual and inconsistently formatted; semantically similar features can be tagged with different keys or values across regions; and temporal drift (stale or rapidly edited entries) leads to inconsistency between adjacent tiles. Ambiguities arise from overlapping tags (`amenity` vs. `tourism`), locally invented values, and mixed geometry types for the same real-world entity. These characteristics constrain downstream classification and further reduce the comparability of results across regions, regardless of the chosen extraction pipeline.

3.2 Classification of Points of Interest

Rule-based classification assigns each retrieved PoI a stable *(category, subcategory)* label that drives filtering, styling, and LoD priorities (Section 3.3). The approach maps raw OSM tags to internal landmark categories through predefined dictionaries and key-specific heuristics, following established landmark extraction practice [Elias and Brenner, 2005]. Its advantages are simplicity, transparency, and predictable runtime.

Data Quality and Robustness. Given the issues outlined in Section 3.1 (uneven completeness, heterogeneous tagging, variable accuracy), classification must tolerate missing names, ambiguous or region-specific tags, and mixed geometries [Al-Bakri, 2015]. The rule-based approach is therefore framed as a conservative mapping from available tags to a stable *(category, subcategory)* label, with explicit documentation of unclassified cases instead of inferring absent semantics. Heuristic fallbacks such as name-based matching provide resilience in practice [Elias and Brenner, 2005], and recent work indicates potential for user-aware adaptation [Zhu et al., 2022].

Taxonomy (UI-Facing Labels). The internal taxonomy organizes landmarks into the two-level *category/subcategory* hierarchy aligned with common OSM tags. The taxonomy is UI-facing while preserving original tag sets for traceability. Typical mappings include *amenity=restaurant → Food & Drink/Restaurant, tourism=viewpoint → Tourism/View*

Point, natural=peak → *Nature/Peak*. The instantiated mapping list is given in Section 4.3 on page 23.

Classification Overview. In practice, classification proceeds in three steps: (i) a deterministic rule lookup over prioritized OSM keys assigns *(category, subcategory)*; (ii) when tags are inconclusive, limited name-based heuristics provide a secondary signal; (iii) records with missing geometry or no informative tags are discarded with an explicit reason, and unclassified cases are logged for later rule refinement.

3.3 Spatial Indexing via Quadtrees and Delivery

Spatial indexing structures are fundamental to rendering efficiency in interactive geographic applications. Quadtrees provide a recursive spatial subdivision that aligns naturally with zoom hierarchies and tile-based rendering [Gaede and Günther, 1998]. In a quadtree, each region is subdivided into four subregions, enabling logarithmic-time lookup for spatially constrained PoIs and supporting multi-scale landmark filtering. A conceptual structure with three representative levels and its subdivision is shown in Figure 3.2. This model supports dynamic LoD rendering, where only PoIs relevant to the current view and zoom level are requested and visualized. In doing so, the index operationalizes scale-dependent filtering mentioned in Chapter 2 as a prerequisite for managing cognitive load in landmark-dense views.

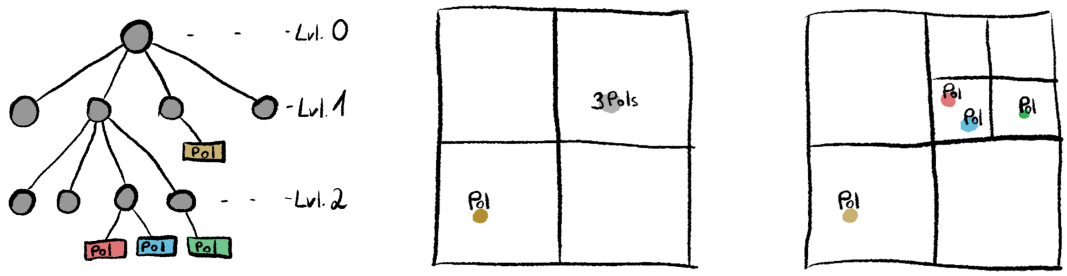

(a) Quadtree tree structure depicting three levels.

(b) Top-down view at level 1; PoIs grouped in parent nodes.

(c) One level deeper at level 2, with individual leaf nodes.

Figure 3.2: Conceptual quadtree structure across levels.

Prior work on continuous LoD for terrain shows that hierarchical control of detail sustains interactive frame rates [Lindstrom et al., 1996]; in landmark visualization, the same hierarchy governs semantic aggregation and refinement. As semantic relevance and user familiarity influence utility [Zhu et al., 2022], the spatial partition is paired with a project-specific, stable category taxonomy and an adaptive delivery path to balance density, clarity, and task needs to ensure efficient filtering and clear visualization at multiple scales.

Structure and Parameters. The quadtree is defined over the project Bounding Box (BBox) with configurable *node capacity* and a fixed *maximum depth* of 5. A BBox is an axis-aligned area defined by two longitudes and two latitudes [OpenStreetMap Wiki, 2024]. Node capacity controls when a cell subdivides, while max depth bounds memory and the smallest spatial grain. This fixed depth is chosen to balance rendering performance and spatial detail, preventing overfragmentation. Each leaf stores compact records (position, category, subcategory, attributes) decoupled from any specific rendering representation. This supports efficient spatial queries decoupled from rendering specifics. Internal parent nodes retain lightweight summaries (representative position and a descendant count) to support clustered rendering without enumerating all items.

Queries. Viewport queries traverse from the root to nodes whose bounds intersect the current viewport. Detail responses enumerate leaf items, while coarser responses return internal-node summaries. The same pruning logic supports local circle queries by rejecting nodes whose bounds cannot intersect the circle before testing distances. Typical forms include *window range* (axis-aligned rectangle), *radius* (within r of a point), and *tile fetch* (by level and tile index).

View-Dependent LOD Strategy. LoD selection is based on camera state and viewport. A continuous "effective zoom" blends the map's discrete zoom with a geometry-derived measure based on camera altitude above ground. Pitch controls the blend to stabilize this behavior under tilt and in mountainous terrain. The continuous value is then mapped to a discrete quadtree depth that determines which level of the quadtree to retrieve the nodes (clusters and details) from.

Data Delivery. Data delivery uses a quadtree index over the point set: for each view, a depth is chosen, only nodes intersecting the viewport are requested, and their PoI records are fetched on demand. Coarser depths provide lightweight overviews; zooming refines by resolving parents to children. A narrow prefetch ring around the viewport reduces boundary latency during panning, and a small in-memory cache smooths short reversals. Terrain mesh detail is handled by the rendering framework. This arrangement provides scale-proportional density, predictable bandwidth use, and a clear separation between access, selection, and visualization. This architecture optimizes performance and ensures a smooth user experience.

3.4 Visualization and Interaction Design

The visualization component translates classified, indexed PoIs into user-facing representations. It is organized as parallel, user-switchable modes that share the same data flow and LoD selection introduced earlier. This parallel design supports task-appropriate choices and structured comparison while keeping interaction behavior consistent across modes.

Goals and Principles. Design goals are: *flexibility* (multiple representational strategies rather than a single style), *scalability* (operation over thousands of PoIs via quadtree filtering and progressive delivery), *clarity* (legible category distinctions with minimal clutter), and *human-centered usability* (predictable lightweight interactions that aid exploration). All modes consume the same quadtree subsets at the same LoD depth; differences stem from symbolization, not from data management or interaction logic.

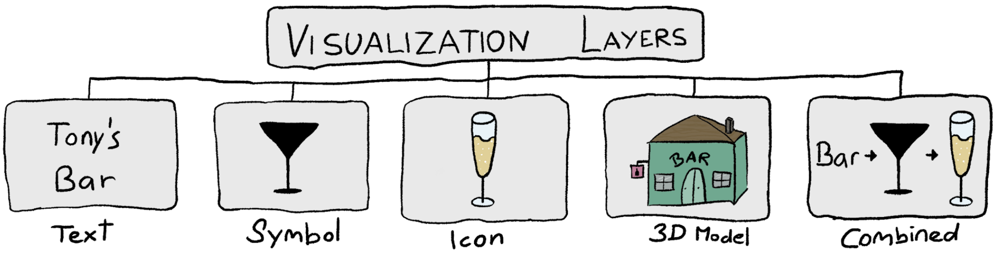

Figure 3.3: Concept of the visualization layers (heatmap not shown).

Visualization Layers (User-Switchable Modes). The system exposes six complementary modes (see Figure 3.3) that share the same data subsets and LoD depth; changing the mode alters only symbolization. In the *text* mode, labels are anchored at landmark positions to provide a transparent semantic baseline; when the view is coarse, clusters appear as grammatically correct count labels placed at representative positions. The *symbol* mode replaces text with abstract 2D symbols that encode the subcategory with low visual complexity, using group symbols for clusters and individual symbols for details. The *icon* mode employs pictorial 2D artwork for rapid category recognition, again switching between group and item assets as the depth changes. The *3D model* mode adds depth cues via generic glTF models that represent subcategories; to balance salience and performance, models are restricted to near ranges while mid/far views fall back to lightweight 2D imagery. The *combined* layer visualizes different LoA by mixing the existing text, icon, and symbol layers depending on distance to the camera. Finally, a complementary *heatmap* presents continuous density overviews for exploration and hotspot finding, with different colors representing each category.

Asset Resolution and Labeling. A unified asset registry maps (*category*, *subcategory*) to filenames for symbols, icons, and 3D models, including cluster variants. Missing subcategory assets fall back to category defaults; absent pictorials still render via symbol or text. Detail features use the `name` attribute when available; clusters display a grammatically correct count label at the node's representative position. Properties (`category`, `subcat`, `label`, `count`) are attached uniformly across modes.

User Interface and Filtering Interaction. To restrict visible PoIs to manageable data subsets, a UI enabling users to select from multiple categories and subcategories based on the hierarchical taxonomy from classification. Selecting one or more categories results in displaying only the corresponding PoIs, effectively reducing clutter and saving cognitive and computational resources. The filtering integrates tightly with the quadtree-based spatial index and LoD strategy, further optimizing performance by selectively loading and rendering data relevant to the task and current view.

Details-on-Demand. Details-on-Demand allows users to access additional important information about PoIs without overwhelming the interface with excessive details. When users hover over a PoI, a lightweight popup appears displaying essential attributes such as the name and category for individual items, or counts for clustered groups.

By clicking on a PoI, the item is pinned to remain visible during navigation, enabling users to maintain focus on relevant details. Pinned items reveal richer attribute information sourced from OSM, conveniently paginated for easy browsing. Notably, the pop-ups and pinned panels consistently present the same set of attributes regardless of the visualization mode, ensuring a uniform user experience across all views.

CHAPTER 4

Implementation

This chapter describes the implementation of the prototype in a web-native environment. Visualization is organized into several comparable layers within the rendering pipeline: text labels, icons, symbols, 3D models, and a categorized heatmap. A combined 2D demonstrator illustrates distance-dependent abstraction without replacing the main visualization layers. The implementation thus serves as both a technical realization and a concrete test bed for investigating the conceptual gaps identified in Chapter 2, specifically in rural terrain environments.

Subsequent sections document the OSM attributes consumed, filtering logic, symbolisation, and interaction hooks. Shared infrastructure, like tile-scoped spatial indexing, event-driven recomputation on `moveend`/`zoomend`, and consistent category theming, provides predictable behavior across layers while allowing composition for analysis and evaluation.

4.1 Libraries and Frameworks

The implementation adopts a web-native stack that balances interactive performance with precise control over PoI overlays and landmark styling. This stack includes Mapbox GL JS as the map-centric foundation offering terrain and imagery rendering, Three.js and Threebox for 3D scene construction and synchronization with the geographic camera, and TypeScript alongside a minimal bundling pipeline to ensure reliability and repeatable builds.

Mapbox GL JS (`mapbox-gl` 3.9.2). Vector-tile renderer with integrated terrain and raster imagery, used for basemap display, camera control, and data-driven styling of overlays [Mapbox, 2025].

Three.js (`three` 0.170.0). General-purpose WebGL framework used to render 3D landmark models and custom scene elements [Three.js, 2025].

Threebox (`threebox-plugin` 2.2.7). Bridge between Mapbox GL JS and Three.js that keeps the 3D scene aligned with the map view, enabling custom layers with georeferenced models [jscastro76, 2025].

TypeScript (`typescript` 5.7.2) with typings `@types/mapbox-gl`, `@types/geojson`, `@types/three`. Adds static types for mapping and rendering APIs to reduce integration errors [Microsoft, 2025].

Build toolchain. `webpack`, `webpack-cli`, `webpack-dev-server`, `ts-loader`. Handles bundling, local development, and TypeScript transpilation for consistent builds [Webpack, 2025].

This configuration follows the rationale established in the background chapter: compared to globe-oriented engines or layer-centric renderers, a map-centric workflow simplifies PoI integration and interaction, while preserving headroom for custom 3D landmark rendering.

4.2 Offline OSM Snapshot (Data Acquisition and Preparation)

Access to OSM data is performed offline and results are cached locally, ensuring reproducibility and predictable throughput. The Austrian project BBox is subdivided into tiles of $0.5° \times 0.5°$ (tile step $\Delta = 0.5°$). For each base tile $B = [\varphi, \varphi + \Delta] \times [\lambda, \lambda + \Delta]$, the Overpass query extent is expanded by a small border margin $\varepsilon = 0.001°$ to avoid truncating features near tile edges.

Acquisition and cache. Parameterized Overpass queries retrieve all features F for later classification and rendering. In the snapshot used for this thesis, the pipeline parsed exactly 6,064,928 features across the area of interest. Requests apply capped exponential backoff on rate-limit responses and timeouts, and failed calls are retried up to $k = 3$. Results are stored as raw OSM XML files in an on-disk cache, keyed by $\langle B, F \rangle$ (base tile B and the feature-selection signature F), preventing repeated identical requests.

Parsing and normalization. XML tiles are parsed to extract point-like PoIs and their attributes (identifier, coordinates, OSM tags). Records lacking valid geometry are discarded. Tags are normalized into a stable internal schema, which supports the two-level category/subcategory taxonomy detailed in Section 3.2.

Deduplication across overlaps. Because the tile overlap is intentional, duplicate elements caused by the margin ε are removed downstream based on their OSM IDs. For geometry fragments that are split across neighboring tiles, the canonical instance is determined by selecting the one whose centroid lies within the non-overlapping base extent of tile B; ties are resolved using the lexicographic order of (φ, λ).

Provenance and timestamping. Each snapshot run is labeled with a UTC timestamp and accompanied by a manifest recording the BBox, Δ, ε, the feature F, the Overpass endpoint, and the commit hash of the query templates. This metadata supports reproducible comparisons between different snapshot versions and is consistent with the rationale described in Section 3.1.

Hand-off to downstream stages. Each element is serialized as `{id, longitude, latitude, category, subcategory, name?, tags}`. Aggregate statistics, such as the total number of parsed PoIs and counts per category and tile, are collected to validate coverage and inform performance tuning for subsequent stages. These preprocessed and labeled PoIs are emitted as compact records, ready for spatial indexing (Section 4.4) and runtime delivery (Section 4.6), avoiding further network calls during interaction.

4.3 Classification of Points of Interest

The classification process operates on entities extracted from the offline OSM snapshot, each of which provides a unique identifier, geographic coordinates, and a set of normalized tags. For non-point geometries that represent landmarks with point-like semantics (e.g., buildings, natural features), centroid coordinates are computed and stored. To preserve traceability, the original geometry and tag set are retained alongside the normalized attributes.

OSM tag and value	**Category**	**Subcategory**
`amenity=restaurant`	Food & Drink	Restaurant
`amenity=drinking_water`	Food & Drink	Drinking Water
`tourism=viewpoint`	Tourism	Viewpoint
`tourism=alpine_hut`	Accommodation	Hut
`natural=peak`	Nature	Peak
`natural=waterfall`	Nature	Waterfall

Table 4.1: Representative tag-to-category mappings.

Deterministic classification order. Classification follows a strict, deterministic precedence across high-level OSM keys to guarantee reproducible results. The system evaluates keys in a fixed priority sequence:

$$\texttt{amenity} \triangleright \texttt{tourism} \triangleright \texttt{shop} \triangleright \texttt{place} \triangleright \texttt{emergency} \triangleright \texttt{natural} \triangleright \texttt{...}$$

The first key present in this order determines the classification path. For that key, its value is compared against a predefined mapping dictionary. If multiple values are present for the same key, the lexicographically smallest one is selected to ensure deterministic behavior across runs. If a dictionary entry matches exactly, the feature is assigned a two-level label consisting of (category, subcategory). Representative mappings are shown in Table 4.1.

Key-specific wildcard rules absorb long-tail values without requiring exhaustive lists. For example, `shop=*` is collapsed into *Shop / Misc Shops*, while `natural=*` is mapped to *Nature / Misc Nature*. This guarantees that even less common values are consistently categorized according to the same hierarchy, without introducing non-deterministic fallbacks.

Fallback heuristics. If no dictionary rule yields a classification, the system applies a restricted set of substring checks over the attributes name and operator. This heuristic layer provides a secondary signal to improve coverage, for example:

Accommodation / Hut: `hut`, `cabin`; *Nature / Peak*: `gipfel`; *Nature / Waterfall*: `waterfall`, `falls`, `cascad`.

Heuristics are triggered strictly when rule-based lookup fails and never override existing exact matches. This guarantees reproducibility and prevents non-deterministic reassignments.

Output. After classification, the system applies strict validation rules to ensure data quality: Records without valid coordinates are discarded with the reason "missing geometry". Records without informative tags after normalization are discarded with the reason "no tags". Features that remain unclassifiable after heuristic checks are logged line by line into `unclassified-pois.json` (JSONL format). Each entry contains a minimal identifier set (ID, longitude, latitude, tags) together with the reason unclassified. This log supports later refinement of classification rules and ensures reproducibility of the process.

Successfully classified PoIs are serialized into structured records of the form id, lon, lat, category, subcategory, name?, tags, where name is included only if available. The output of this stage is organized into per-subcategory groups, which serve as direct input to the spatial indexing step described in Section 4.4.

4.4 Spatial Index Construction

The spatial index is defined over the Austrian project area $A = [\varphi_{\min}, \varphi_{\max}] \times [\lambda_{\min}, \lambda_{\max}]$ in WGS84 (latitude φ, longitude λ). The area is partitioned into a fixed 4×4 grid of *base tiles* by uniform steps

$$\Delta\varphi = \frac{\varphi_{\max} - \varphi_{\min}}{4}, \qquad \Delta\lambda = \frac{\lambda_{\max} - \lambda_{\min}}{4}.$$

Each base tile is addressed by integer indices $(r, c) \in \{0, 1, 2, 3\}^2$ with bounds

$$T_{r,c} = [\varphi_{\min} + r\,\Delta\varphi,\ \varphi_{\min} + (r+1)\,\Delta\varphi] \times [\lambda_{\min} + c\,\Delta\lambda,\ \lambda_{\min} + (c+1)\,\Delta\lambda].$$

For each base tile $T_{r,c}$ and every $(\mathit{category}, \mathit{subcategory})$, an independent quadtree is constructed. Level 0 spans the bounds of $T_{r,c}$, and each deeper level doubles the spatial resolution per axis. This results in $16 \times N_{\text{subcat}}$ trees for any single subcategory.

The fixed 4×4 base partition ensures compact per-tree files, which is critical for efficient runtime serving, and maintains high spatial locality, so a typical viewport overlaps only a few tiles. Independent tree files can be loaded in parallel and updated incrementally when a snapshot changes only part of the area A. Maintaining a separate tree per $(\mathit{category}, \mathit{subcategory})$ preserves semantic homogeneity in internal summaries and simplifies symbolization. Anchoring level 0 at the base-tile bounds ensures that power-of-two subdivisions remain aligned within each tile regardless of the aspect ratio of A, which stabilises the mapping from effective zoom to depth discussed in Section 4.5.

Aggregation strategy and cluster anchors. Each node summarises the PoIs within its quadtree cell by a single aggregate. Internal nodes represent all of their descendants, while leaf nodes represent only their own members. Two statistics are stored: *poiCount* (the number of contained items) and *averagePosition* (the arithmetic mean of descendant coordinates in (λ, φ)). Let $S(n)$ denote the multiset of PoIs in node n. Then

$$\mathit{poiCount}(n) = |S(n)|, \qquad \mathit{averagePosition}(n) = \frac{1}{|S(n)|} \sum_{p \in S(n)} (\lambda_p, \varphi_p).$$

For implementation efficiency, the *averagePosition* of internal nodes is computed as the count-weighted mean of the children's *averagePosition*. At coarser depths, clusters are anchored at *averagePosition* and labelled with *poiCount*. Because aggregation is performed in object space, cluster identities remain stable during panning and are invariant to the current field of view. For the Austrian extent and chosen tile sizes, the planar average in geographic coordinates provides a sufficiently accurate approximation. Since cells are small and far from the antimeridian, geodesic corrections are negligible at the applied depths.

4.5 View-Dependent LoD and Traversal

Effective zoom and depth selection. Altitude above ground is estimated by sampling the terrain at the four viewport corners and the centre, and subtracting these values from the camera altitude. Using this altitude and the vertical field of view, the visible span is computed as:

$$\text{span} = 2 \cdot \text{altitude}_{\text{AGL}} \cdot \tan(\text{fov}/2).$$

The geometry-based zoom is defined as a base-2 logarithmic ratio of Earth circumference $C_\oplus$ to span:

$$\text{geometryZoom} = \log_2\left(\frac{C_\oplus}{\text{span}}\right).$$

Pitch is restricted to the interval $[30°, 60°]$. A pitch-aware weight t increases linearly from 0 at $30°$ to 1 at $60°$. The effective zoom is computed as a weighted combination of the map's nominal zoom and the geometry-based zoom:

$$\text{effectiveZoom} = (1 - t) \cdot \text{mapZoom} + t \cdot \text{geometryZoom}.$$

The effective zoom is then mapped to a discrete quadtree depth in the range $1 \ldots 6$ using fixed thresholds:

$$\leq 8 \mapsto 1, \quad \leq 9 \mapsto 2, \quad \leq 12 \mapsto 3, \quad \leq 14 \mapsto 4, \quad \leq 15 \mapsto 5, \quad \text{else } 6.$$

Depth 1 represents the coarsest level used at runtime. Transitions between depths are strictly discrete, and this module does not apply temporal smoothing or opacity blending.

Traversal and emission. For the selected (*category*, *subcategory*) and depth level, all nodes belonging to that level are collected. To prevent gaps in sparse regions, leaves from higher levels are also included. Internal nodes emit PoIs of `type='cluster'`, anchored at their *averagePosition*. Leaf nodes emit PoIs of `type='detail'` containing the full PoI record. As illustrated in Figure 4.1, parent nodes appear as clusters at overview scale (a), subdivide into smaller clusters upon zooming (b), and reveal individual PoIs at leaf level (c).

(a) Overview. (b) Zoomed. (c) Detail.

Figure 4.1: Applied quadtree structure.

Detail on demand and local counts. Clusters can be expanded locally without modifying the global quadtree depth, enabling details-on-demand functionality. Local counts are computed via circle queries, which prune branches using BBox–circle intersection before calculating precise Haversine distances for the remaining leaf nodes. This method similarly facilitates summarization per category and per subcategory.

Output schema. The traversal produces GeoJSON Point features with a consistent property schema used for styling: type `'cluster'` or `'detail'`, `category`, `subcat`, `iconName`, `symbolName`, `modelName`, `icon3DName`, and, where applicable, `count` and `label`. Switching layers affects only the symbolization, while spatial granularity is governed solely by the chosen depth level.

4.6 Runtime Data Service

The runtime data service reuses the fixed 4×4 base-tile grid established for indexing in Section 4.4. Each chunk corresponds to one base tile and is addressed via stable integer indices $(r, c) \in \{0, 1, 2, 3\}^2$, arranged in row-major order in both filenames and identifiers. Requests are issued only for chunks whose spatial extent intersects the current viewport; all others remain inactive. Each payload includes the serialized quadtree for each $(category, subcategory)$ within the chunk, preceded by a compact header containing tile bounds, format version, and a provenance manifest pointer.

Priorities, prefetch, and cancellation. Chunk visibility drives request priority, with chunks overlapping the viewport enqueued first. To buffer minor camera movements, the immediate one-ring of neighbouring chunks is prefetched at lower priority. Should the viewport change before a response arrives, stale requests for chunks no longer intersecting the viewport are cancelled, duplicates merged, and only the most recent request retained for each (r, c).

Caching and resilience. A bounded in-memory least-recently-used (LRU) cache stores recently fetched chunk payloads, enabling high-speed retrieval on revisits. Missing or malformed responses result in lightweight placeholders that allow rendering to continue; as valid data arrive, clusters or details are refined dynamically. Basic validations on content type and maximum payload size are enforced, and diagnostic per-request timings are logged.

Output. From the renderer's point of view, the service returns either a populated chunk containing the header and per-category trees or an empty placeholder. This delivery mechanism remains agnostic to the LoD; decisions concerning depth selection and cluster versus detail materialization are managed exclusively by Section 4.5.

4.7 Multi-Modal Visualization and Levels of Abstraction

Parallel, user-switchable visualization modes enable direct comparison of abstract, textual, symbolic, and detailed 3D representations while operating at a consistent LoD (Section 4.5). The prototype supports six modes: Text, Symbol, Icon, Combined LoA, Heatmap, and 3D model. All modes draw from the same GeoJSON point source. Each feature carries uniform properties, ensuring consistency across modes. Asset references (icons and models) follow a hierarchical fallback from subcategory to category to global defaults, and downloaded resources are cached to prevent redundant requests and minimize loading times.

2D Visualization Modes

Text Layer. Labels are placed directly at the location of each PoI, providing an explicit semantic cue. Text elements remain legible across zoom levels and incur minimal rendering overhead, serving as the foundational representation (Figure 4.2a).

Icon Layer. Pictorial markers evoke real-world objects to enhance immediate recognition without compromising spatial accuracy. Icons balance expressive detail and clarity (Figure 4.2b), making them suitable for mid-range identification tasks.

Symbol Layer. Abstract symbols represent categories with simple shapes, reducing map clutter while preserving precise geographic alignment. By limiting visual complexity, this mode supports rapid overview and comparison (Figure 4.2c).

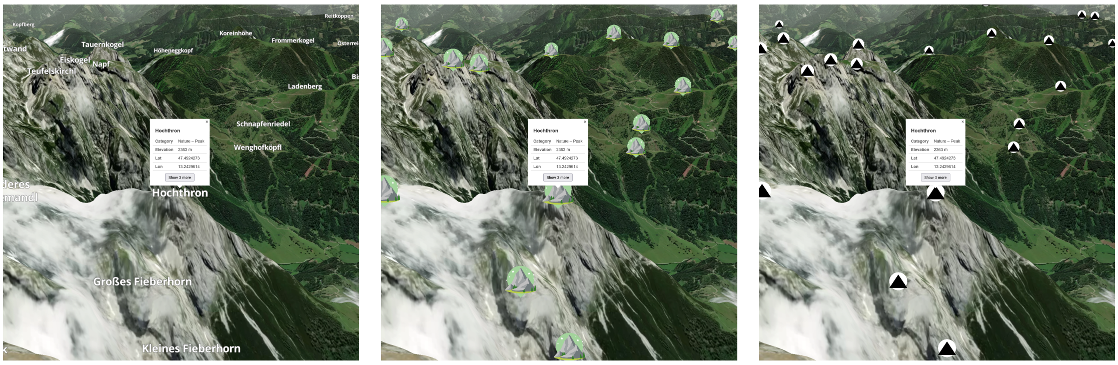

(a) Text label mode. (b) Icon mode. (c) Symbol mode.

Figure 4.2: Hover interaction of the same PoI in different visualization modes.

Combined Layer with different Levels of Abstraction

This two-dimensional layer demonstrates distance-dependent abstraction by blending three representations (symbols for far range, icons for mid range, and text labels for close range) from the same GeoJSON source (Section 3.4). Near, mid, and far thresholds are loaded from the application state and injected into Mapbox filter expressions on three dedicated sublayers. Each sublayer's filter uses the precomputed dist property to select appropriate features. A quadtree-based tile limiter restricts per-frame computations to visible tiles and their neighbors, keeping interactions performant. This style-driven approach ensures that switching abstraction bands requires only filter re-evaluation instead of manual re-rendering, preserving a coherent visual hierarchy as the user navigates. The result is a deterministic abstraction schedule that clarifies the overview, supports mid-scale search, and maintains legible close-range identification.

3D Visualization Mode

To balance salience and performance, the 3D layer applies a distance-dependent staging policy integrated with the quadtree LoD system described in Section 4.5. This policy leverages the quadtree's cluster/detail node distinction so that computational resources are allocated efficiently while maintaining visual clarity for landmark recognition and allowing allocation of geometric detail only when it improves recognition [Kapaj et al., 2022, Kapaj et al., 2024, Carbonell Carrera et al., 2017, Çöltekin et al., 2018, Smallman and Cook, 2011].

Figure 4.3: Activated 3D layer visualizing a generic hut at close range.

Architecture and Integration. The 3D rendering layer is implemented as a custom Mapbox layer utilizing Threebox to provide seamless interoperability with existing 2D cartographic layers. This architecture enables efficient coordinate transformation between geographic and rendering coordinate systems while maintaining consistent camera synchronization across visualization modes.

Distance-Driven Model Staging. To balance performance and visual fidelity, PoIs first appear as 2D impostors, which are flat snapshots of their 3D models with transparent backgrounds. When the camera comes within 200m, each impostor is replaced by a full glTF model clamped to the terrain. A small margin around this threshold prevents rapid toggling between representations. This staging uses the quadtree LoD hierarchy. Cluster nodes always remain as 2D impostors, while detail nodes switch to full 3D models when they enter the near band. By combining staging logic with quadtree traversal, the system instantiates only the models needed for the current view, thereby reducing draw calls and CPU load without sacrificing spatial detail where it matters.

Asset Management and Model Preparation. Assets are centrally managed via a registry that maps each category to its glTF URL and fallback placeholder. During initialization, the loader module fetches, parses, and caches these models, ensuring that only needed assets are requested and that cached objects are ready for instantiation, which minimizes both load latency and memory overhead. All models use the glTF format for its compact size and fast parsing in web environments. To maintain consistent visual prominence, each model's BBox is scaled to a uniform target size upon loading. This combined pipeline unifies asset resolution, efficient delivery, and runtime scaling without requiring separate handling for format conversion or sizing.

Terrain Integration and Memory Management. 3D models are positioned using terrain elevation queries, supporting both absolute and relative height placement. A disposal routine cleans up Three.js geometries and materials when models are removed, preventing memory leaks and ensuring smooth performance during extended use in rural environments.

Categorized Heatmap Layer

Mapbox does not natively support terrain-draped heatmaps in 3D scenes, so this layer renders circular kernels in screen space, with only each kernel's center point projected onto the terrain surface (Figure 4.4). On steep slopes or at high zoom levels, the planar kernels detach visually from the ground, causing slight positional ambiguity along their edges. However, the overall effect provides a useful density overview (Figure 4.4a) while highlighting relief-related limitations (Figure 4.4b).

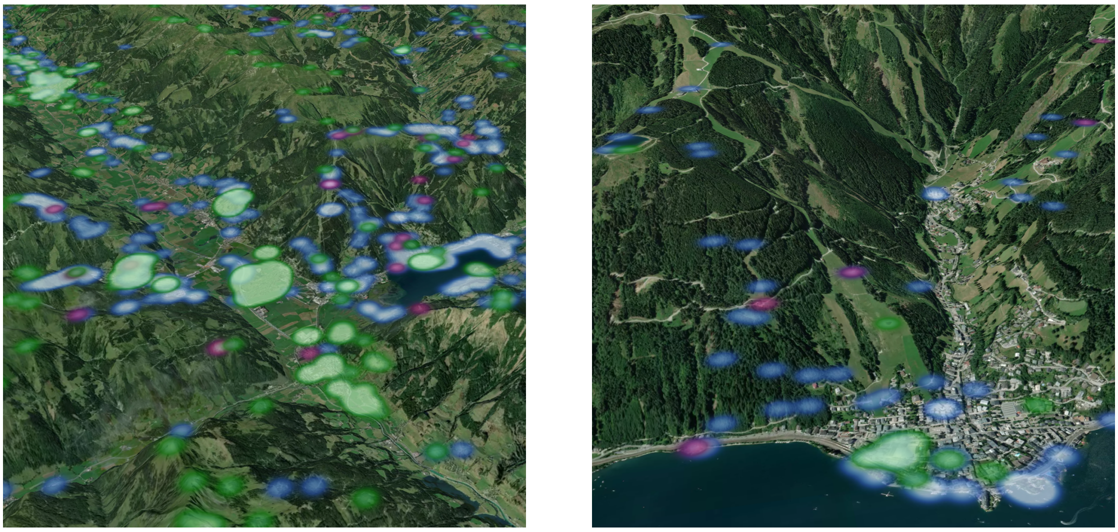

(a) Categorized heatmap at far range. (b) Categorized heatmap at close range.

Figure 4.4: Heatmaps with category-specific colors.

Rather than a single global style, PoIs are aggregated by category and mapped to distinct color themes. Each category uses a configurable kernel radius and opacity ramp computed via a Gaussian weighting function to emphasize hotspots without overwhelming the view. At rendering time, the layer issues a single WebGL draw call per category, leveraging instanced circle geometries for performance.

4.8 User Interface Components and Interactions

Category-based filtering. To aid wayfinding at decision points without overwhelming the user, the interface provides a hierarchical category filter derived from preprocessing and classification (Section 4.2, Section 4.3). Prior work shows that salience peaks with three to five items and declines beyond seven, increasing cognitive load [Cheng et al., 2022]. Going in this direction, users can toggle categories and subcategories to regulate landmark density (Figure 4.5).

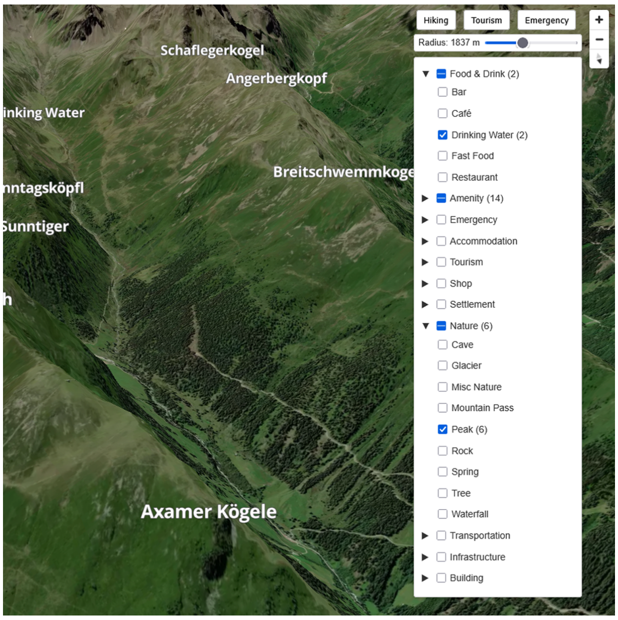

Figure 4.5: Category-based filtering interface to regulate PoI density.

Quick selection buttons. To streamline mode changes for common use cases, the interface provides three preset buttons positioned above the category filter (Figure 4.5). Hiking, tourism, and emergency contexts are mapped to a deterministic set of subcategories. For instance, the hiking preset simultaneously activates peaks, drinking water, benches, and viewpoints. Upon clicking a preset, the visibility state for all associated subcategories is updated and the map view refreshed in a single operation, reducing user effort and cognitive overhead in scenario-driven navigation.

Nearby PoIs and radius control. A configurable geographic radius around the viewport center defines a clipped area of interest on the map surface. As users move the slider, a filled, terrain-conforming circle is drawn on the map to indicate the current search boundary, and the count of PoIs within that boundary updates in real time as the map is panned or zoomed (Figure 4.6). This dynamic display allows users to explore local landmark density and spatial context interactively.

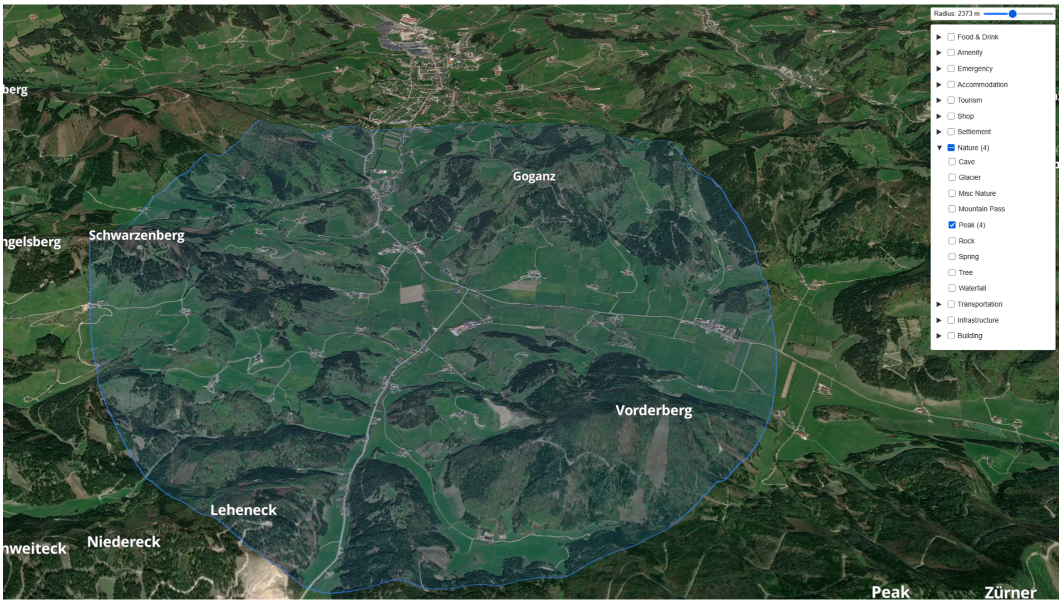

Figure 4.6: A slider adjusts the search radius for nearby PoIs around the center of the viewport, visualized as a circle in the map.

Popups and the information panel. Selecting a PoI opens a persistent information panel displaying its name, category, and relevant OSM tags (see Figure 4.2, page 28). Because OSM entries vary in attribute richness, the panel paginates tag data into manageable sections, and only one panel remains open at a time. Hovering over a PoI reveals a lightweight popup that disappears when the pointer moves away; clicking the popup converts it into the persistent pane. This two-tier interaction is implemented via Mapbox event listeners on feature layers that dynamically generate and inject HTML content based on each PoI's metadata. This pattern supports rapid inspection and focused comparison, aligns with cognitive load management principles through progressive disclosure [Smallman and Cook, 2011, Yesiltepe et al., 2021], and leverages landmarks as anchors to improve wayfinding efficiency [Cliburn et al., 2007, Cheng et al., 2022].

CHAPTER 5

Evaluation and Results

This chapter evaluates the developed prototype in relation to the central research question introduced in Chapter 1. The evaluation adopts a reflective, design-oriented approach: instead of a controlled user study, the analysis examines to what extent the implemented choices address the gaps outlined in Chapter 2 and realize the design principles introduced in Chapter 3.

The evaluation is structured along four analytical dimensions: landmark representation, human-centred design considerations, LoD and delivery, and data integration and classification. These dimensions correspond directly to the implemented visualization layers (text, icons, symbols, selected 3D models, the categorized heatmap, and the combined 2D LoA demonstrator). Table 5.1 provides an overview of the roles of these layers with respect to task suitability and viewing distance.

To ensure methodological rigour, the discussion is guided by explicitly defined criteria and supported through traceable reasoning consistently referenced in Table 5.1. For each dimension, the relevant criteria are briefly stated, the implemented design is mapped against them, and a reflective assessment is provided in the context of rural, terrain-based visualization. The chapter concludes with a synthesis (Section 5.5) and a discussion of limitations and future implications (Section 5.6).

5.1 Landmark Representation

The prototype implements multiple modes of landmark visualization: two-dimensional symbols and icons, textual labels, selected 3D models, and a categorized heatmap. The availability of multiple modes enables within-system comparison, as layers can be toggled and combined to examine how the abstraction level, symbol form, and density affect understanding across scales. The following discussion focuses on realism versus abstraction, symbol modality and legibility, and density with scale-dependent filtering, and is consolidated in Table 5.1.

Realism and Abstraction

Evidence suggests a nuanced role of realism: realistic 3D depictions can aid recognition and route learning, especially for users with lower spatial abilities [Kapaj et al., 2022, Kapaj et al., 2024], and visual richness may direct attention to navigation-relevant features [Nurminen and Sirvio, 2021]. At the same time, excessive detail can compete with task cues and impair effectiveness [Smallman and Cook, 2011, Çöltekin et al., 2018]. Realism is therefore applied selectively. Detailed 3D models are confined to categories with distinctive volumetric form (e.g., huts, trees, churches). Symbols use austere, monochrome geometry to stabilise figure–ground at overview scales. Icons provide a mid-level with recognizable pictorial detail and category-colored background for rapid parsing. This tiered allocation of realism preserves hierarchy: abstraction for scalable overview, pictorial detail supports mid-scale search, and 3D form for close-range recognition. This way, it avoids the pitfalls of naïve realism [Smallman and Cook, 2011, Yesiltepe et al., 2021]. LoA are treated as established means of controlling thematic granularity and stylisation [Semmo et al., 2012, Pasewaldt et al., 2012].

With respect to the research question, this allocation operationalises a concrete design rule for rural 3D terrain: abstraction ensures global legibility, pictorial detail supports category search, and selective 3D benefits close-range comprehension where distinctive volumetric shapes are present. Through this tiered use of representation, the prototype demonstrates how visual choices can be systematically tuned to task and distance to improve spatial exploration and landmark understanding, as consolidated in Table 5.1.

Symbol Types and Legibility

Symbol modality influences both attention and memory. Text labels, set in white with a black outline, prioritise legibility across heterogeneous terrain and shading, and have been shown to support superior recall despite comparable fixation durations [Franke and Schweikart, 2017]. Symbols minimise graphical complexity to convey structure and presence under higher densities, employing austere geometry that signals category with low visual load (Figure 5.1). Icons provide recognizable silhouettes combined with category-colored backgrounds, which accelerate mid scale identification and category parsing (Figure 5.2). Selected 3D models complement these roles at close range, where distinctive volumetric form aids recognition. All artwork is open-source or created by the author using open-source assets.

Landmark Density and Scale-Dependent Filtering

The number of simultaneously visible landmarks is a key determinant of cognitive load. Empirical studies demonstrate that benefits plateau once a modest set of salient cues is available (approximately five at a decision point) and may even decline when additional

(a) Peak. (b) Bar. (c) Bus Station. (d) Café. (e) University.

Figure 5.1: Examples of symbols used in the 2D symbol layer.

(a) Peak. (b) Bar. (c) Bus Station. (d) Café. (e) University.

Figure 5.2: Examples of pictorial icons with category-specific backgrounds.

items are shown [Cheng et al., 2022]. Addressing this, the prototype does not rely on a fixed global cap but instead applies the dynamic LoD pipeline introduced conceptually in Chapter 3 (Section 3.3) and technically realized in Chapter 4 (Section 4.5). This mechanism clusters dense regions and filters PoIs by both scale and category, keeping on-screen density responsive to current view conditions. In addition, the per-layer toggles and category filters defined in the UI design (Chapter 4, Section 4.8) enable interactive adjustment, allowing users to balance information richness with the principle of a "useful many" without introducing overload [Yesiltepe et al., 2021].

Synthesis: Comparative Affordances by Task and Distance

The implemented visualization layers are designed to be complementary rather than redundant, with each addressing a distinct segment of the task–distance space. As consolidated in Table 5.1, their expected strengths, limitations, and typical uses derive both from established findings in the literature and from the conceptual design rules introduced in Chapter 3 and their technical realization in Chapter 4.

The combined two-dimensional layer further operationalizes these findings by applying distance-dependent LoA simultaneously (symbols at far range, icons at mid range, and text at close range), thereby instantiating the task–distance mapping within a single integrated representation. This layered composition reflects the conceptual framework of scale-dependent visualization (Chapter 3, Section 3.3) and demonstrates its feasibility in a web-native prototype (Chapter 4, Section 4.7).

In summary, symbols provide scalable category overviews at far distances; icons accelerate mid-scale category search; labels secure precise identification and instruction at decision

Table 5.1: Comparative affordances of landmark layers in terms of strengths, limitations, and task-distance suitability.

Layer	Primary strengths	Main limitations	Best suited for
Symbols: 2D, abstract, black	• Minimal ink; low clutter • Rapid structure and pattern reading • Robust against basemap variation	• Low semantic specificity • Weaker recognition at decision points	• Far-distance overview • Assessment of category distribution
Icons: 2D, pictorial, category background	• Fast category parsing • recognizable shapes at mid-scale • Good balance of meaning and visual economy	• Competition at high densities • Color interaction with basemap • Risk of over-specific detail [Smallman and Cook, 2011]	• Route preview; mid-zoom planning • Finding relevant categories near route
Text labels: white with black outline	• Highest recall for named landmarks [Franke and Schweikart, 2017] • Unambiguous identification • Strong contrast across different terrains	• Conflicts/occlusion at high density • Slower initial detection than pictorial marks	• Selection and confirmation • Instructions • Decision points [Yesiltepe et al., 2021]
3D models: selected categories	• Close range recognition • Supports recognition of distinctive forms with lower spatial abilities [Kapaj et al., 2022, Kapaj et al., 2024] • Can guide attention [Nurminen and Sirvio, 2021]	• Performance and clutter costs • Limited benefit for some tasks • Potential visual competition [Smallman and Cook, 2011, Çöltekin et al., 2018]	• Local inspection • Confirmation at near-ground or oblique views
Heatmap: continuous density	• Reveals clusters and gradients without symbol clutter • Robust to positional noise • Category-specific color aid scanning	• No discrete reference cues • Normalisation and blending can mislead [Çöltekin et al., 2018, Keil et al., 2020, Yesiltepe et al., 2021]	• Macro overview • Hotspot reconnaissance at far–mid zoom

points; 3D models support close-range recognition of volumetric forms; and the heatmap enables macro-level cluster reconnaissance before switching to discrete layers (Table 5.1).

Reflective Assessment Selectivity in realism, evidence-based label design, and density management via LoD align with perceptual and cognitive findings [Cheng et al., 2022]. The ability to compare modalities in situ is a strength for analysing trade-offs between overview and identification. At the same time, claims about relative effectiveness remain bounded by the absence of controlled user testing in rural 3D contexts.

These design choices closely reflect the conceptual considerations introduced in Chapter 3, notably the scale-dependent filtering via quadtrees (Section 3.3) and the emphasis on flexible, user-switchable modes (Section 3.4). Their feasibility has been demonstrated in the implementation of the traversal and LoD pipeline (Chapter 4, Section 4.5) and the interaction layer (Chapter 4, Section 4.8). Taken together, this grounded design provides an extensible basis for subsequent empirical evaluation in real-world rural terrain settings.

5.2 Human-Centered Design Considerations

Beyond technical feasibility, the effectiveness of landmark-based navigation depends on alignment with perceptual and cognitive constraints. This section evaluates how interaction choices, symbol design, and density management relate to prior findings from perception and cognition research and to the mechanisms specified in the conceptual framework, including scale-dependent filtering and interaction design (Sections 3.3, 3.4), as well as their technical implementation in the LoD pipeline and UI layer (Sections 4.5, 4.7).

Salience, Visibility and Realism

Landmark effectiveness is closely tied to perceptual salience and recognizability. Reviews emphasize that landmarks are most useful when visually distinctive, semantically meaningful, and consistently recognizable, especially at decision points [Yesiltepe et al., 2021]. These conditions are addressed both *between* modalities, through the complementary roles of symbols, icons, labels, and selected 3D models (Section 5.1), and *within* modalities, through subcategory differentiation (symbols remain minimal yet distinct, icons combine recognizable silhouettes with category-colored backgrounds). Visibility cues are maintained by the view-dependent selection and filtering pipeline introduced in the concept (Sections 3.3, 3.4) and realized in the implementation (Sections 4.5, 4.7). Consistent with evidence that excessive realism can compete with task cues [Smallman and Cook, 2011, Çöltekin et al., 2018], realism is applied selectively: 3D models are confined to categories whose volumetric form facilitates close-range recognition and support users with lower spatial abilities who benefit from richer cues [Kapaj et al., 2022, Kapaj et al., 2024], while schematic modalities maintain scalability for overview and mid-scale search (Table 5.1).

Cognitive Load and Information Density

Cognitive load rises with the number and heterogeneity of simultaneously visible items; benefits plateau around a modest set of salient cues [Cheng et al., 2022]. The system manages cognitive load using the scale-dependent LoD and clustering pipeline specified in Chapter 3 and technically implemented through tile-scoped indexing in Chapter 4 (Section 4.4). At low zoom levels, dense regions aggregate; as scale increases, clusters resolve into individual items, yielding a controlled transition from overview to detail. Category filters and per-layer toggles (Sections 4.7, 4.8) further constrain on-screen complexity to task-relevant classes.

Interaction Design

Interaction externalises memory and reduces context switching (Section 3.4). Hover tooltips offer succinct information disclosure; clicks pin tooltips until dismissed, with two simultaneous pins able to be maintained across pan and zoom within a session. Quick-selection presets activate task-oriented category sets with a single action, lowering interaction cost for common contexts. These mechanisms provide stable references and controlled information flow.

Reflective Assessment The human-centered aspects of the prototype are theoretically well-motivated: salience is addressed through between- and within-mode differentiation; cognitive load is managed via scale-dependent density control; and interaction design supports persistence and task-focused filtering. While effectiveness has not been validated in a controlled study, the design decisions are literature-grounded and explicit, and Table 5.1 formulates testable hypotheses linking modalities to task–distance fit.

5.3 Level of Detail and Delivery

After establishing representational options and human-centred criteria, attention now turns to the mechanisms that keep information density tractable and interaction fluid. Whereas Section 5.1 addressed *how* landmarks are depicted, this section focuses on *how many* are depicted and *how quickly* they are accessed. In this prototype, LoD is applied to PoIs; mesh-level terrain LoD is managed by the underlying framework and remains out of scope.

In relation to the research question, LoD functions here as the primary clutter-management device for rural maps: density is moderated by scale rather than through global caps, ensuring that overview, category search, and close-range confirmation remain legible.

Scale-Aware Aggregation via Quadtree

On-screen density is controlled by a quadtree over the landmark set, introduced in Chapter 3 and realised in Section 4.4. For a given view, the controller selects a depth level and activates only nodes intersecting the viewport at that level; coarse levels favour overview, intermediate depths expose mid-scale search, and fine levels are used for identification. Category filters and local clustering further narrow the number of active items. Two key properties follow: density grows proportionally with scale rather than being limited by a global cap, and eligibility is decoupled from symbolisation.

Hierarchical Indexing for Access

The same hierarchy enables view-dependent retrieval. Each node stores a BBox and compact metadata, ensuring that only nodes intersecting the viewport are requested under pan and zoom, consistent with established spatial access methods [Gaede and Günther, 1998]. Unifying selection and retrieval in a single data structure mirrors tile-based web mapping practices and supports the responsiveness required for mid-scale icon search and near-scale label/3D confirmation as outlined in Table 5.1.

Delivery and Conservative Prefetching

Landmark data are chunked by quadtree node and streamed on demand. To reduce boundary latency during panning, the system performs opportunistic prefetching of the immediate ring of neighbouring tiles; this bounded ring helps maintain predictable bandwidth in rural contexts. Such anticipatory staging aligns with prefetching strategies demonstrated to improve interactive exploration [Battle et al., 2016].

Reflective Assessment The approach emphasizes scale-proportional density, couples access and selection in a single hierarchy, and employs a simple, extensible prefetching strategy. These design choices align with classical work on spatial indexing [Gaede and Günther, 1998] and findings on latency reductions through staging [Battle et al., 2016].

A current limitation is the absence of systematic benchmarks, such as fetch latency under varying bandwidth conditions, cache hit ratios, and frame times at peak densities. Future work should focus on quantifying these metrics and exploring adaptive policies, for example, widening the prefetch ring during rapid pans, alongside cache eviction tuned to rural sparsity patterns.

5.4 Data Integration and Classification

This section evaluates how OSM sourcing, cleaning, and rule-based classification support the scale-aware landmark pipeline described in Chapters 3 (see Sections 3.1 and 3.2). The focus lies on semantic completeness, category mapping, and rural suitability, as these choices underpin the task–distance roles of the layers in Table 5.1: reliable names are essential for text labels, and visual representations are based on consistent categories.

Automatic Landmark Selection and Classification

A foundational contribution to landmark generation was presented by [Elias and Brenner, 2005], who proposed a two-stage process consisting of attribute-based salience measures and visibility analysis using digital surface models. The prototype adapts the first stage to OSM by filtering tags and mapping results to a predefined category–subcategory hierarchy (Section 3.2). This generalises landmark selection from authoritative sources to volunteered data, while recognising that LoA and depiction are addressed separately (Sections 3.3, 3.4). Where surface models are available, optional viewport screening can be introduced as a lightweight second stage in future work.

Data Quality in Rural Contexts

A known challenge of using OSM is the variable completeness and positional accuracy of features, especially outside urban areas. Reviews of OSM quality emphasize that its suitability depends on application requirements [Al-Bakri, 2015]. Gaps in naming, misclassification, or under-tagging of perceptually salient features limit downstream consistency.

Mitigations used in the prototype prioritise features with strong semantic cues (e.g., named peaks) and apply conservative, rule-based mapping (Section 3.2). Nevertheless, the robustness of category assignment and naming remains contingent on local contribution patterns. In contrast to earlier proposals for standardized landmark categories [Yesiltepe et al., 2021], the prototype must operate with the available tags, which sometimes underrepresent culturally or perceptually salient features.

Reflective Assessment An OSM-based pipeline scales and transfers well, operationalizing established selection principles [Elias and Brenner, 2005], but inherits rural variability [Al-Bakri, 2015] and heterogeneous tag semantics. The affordances presented in Table 5.1 should therefore be interpreted as conditional on minimum data-quality thresholds. Practical refinements include quality checks for `name` and category completeness, simple blacklist/whitelist rules, sampling-based audits, and optional visibility screening where surface models exist. These measures increase the likelihood that symbol-layer assignments reflect the intended task–distance roles in rural 3D settings.

5.5 Synthesis: Addressing the Research Question

Synthesising key aspects of representation, human-centred considerations, LoD management, and data integration reveals that effective landmark visualization in rural 3D terrain necessitates a hybrid, scale-aware approach grounded in well-defined access and interaction principles, as detailed in Sections 3.1–3.4 and 4.2–4.8. This synthesis is followed by a qualitative examination of three distinct rural cases (Figures 5.3–5.5), interpreting observations in relation to previous research and formulating corresponding hypotheses.

Qualitative Inspection: Cases and Visual Evidence

To anchor the synthesis in representative rural contexts, three locations with contrasting characteristics were selected to capture variation in terrain-induced occlusion, feature density, and semantic diversity. These factors jointly influence symbol detectability, scale transitions, and density management. The selected sites include an alpine mountain range, a small rural town, and a low-relief plain.

In the alpine mountain range, steep terrain and frequent occlusion challenge far-to-mid scale transitions and highlight interactions between shaded relief and symbol contrast. The small rural town features moderate relief and a compact street network, resulting in higher semantic density at decision points, enabling evaluation of label conflicts, icon parsing, and 3D confirmation. The low-relief plain offers wide open views with minimal occlusion, emphasizing long-range scanning and category clustering, thus revealing the boundaries between symbol overviews, icon emergence, and label confirmation.

(a) Symbols (mid).

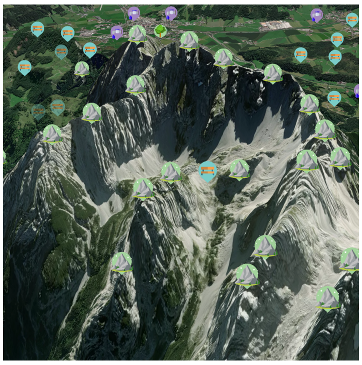

(b) Icons (mid).

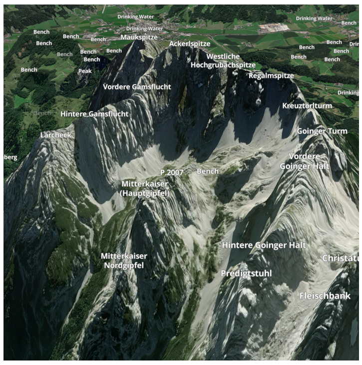

(c) Text labels (mid).

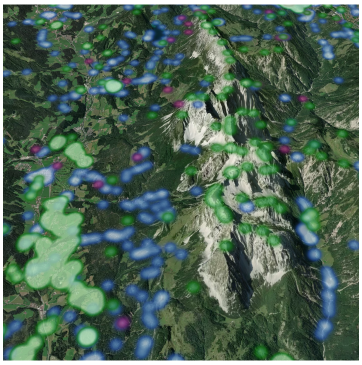

(d) Heatmap (far).

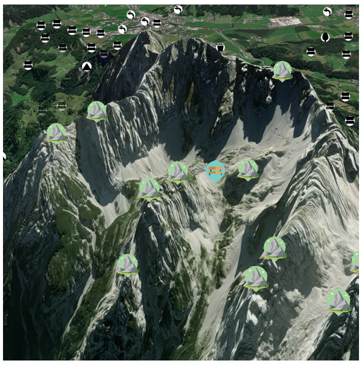

(e) Combined (mid).

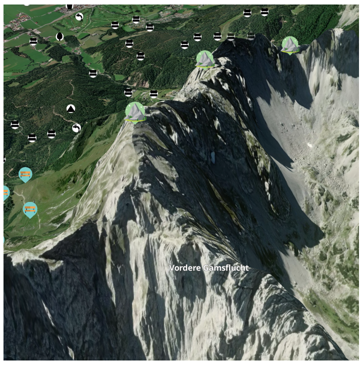

(f) Combined (near).

Figure 5.3: Alpine mountain range.

Procedure. For each of the three selected rural locations, analysis was conducted at three zoom levels aligned with the defined LoA thresholds: far, mid, and near. The evaluation focused on multiple visualization modalities, including discrete PoIs represented as symbols, icons, and text labels, along with the relevant 3D models. Heatmaps were also examined to provide macro-scale overviews of landmark density and distribution. These diverse layers were assessed collectively through distance-based combinations to simulate realistic navigation conditions.

Observations concentrated on key factors influencing landmark visualization effectiveness: occlusion, symbol conflicts, detectability, category parsing, and appropriateness for specific spatial tasks. The assessment was guided by the comparative strengths, limitations, and task-distance suitability detailed in Table 5.1, integrating insights from both the conceptual framework (Chapter 3) and the technical implementation (Chapter 4).

These observations form the empirical foundation for the subsequent synthesis, which addresses the central research question:

How should landmark visualization modes be designed to support spatial exploration and landmark comprehension in rural 3D terrain maps?

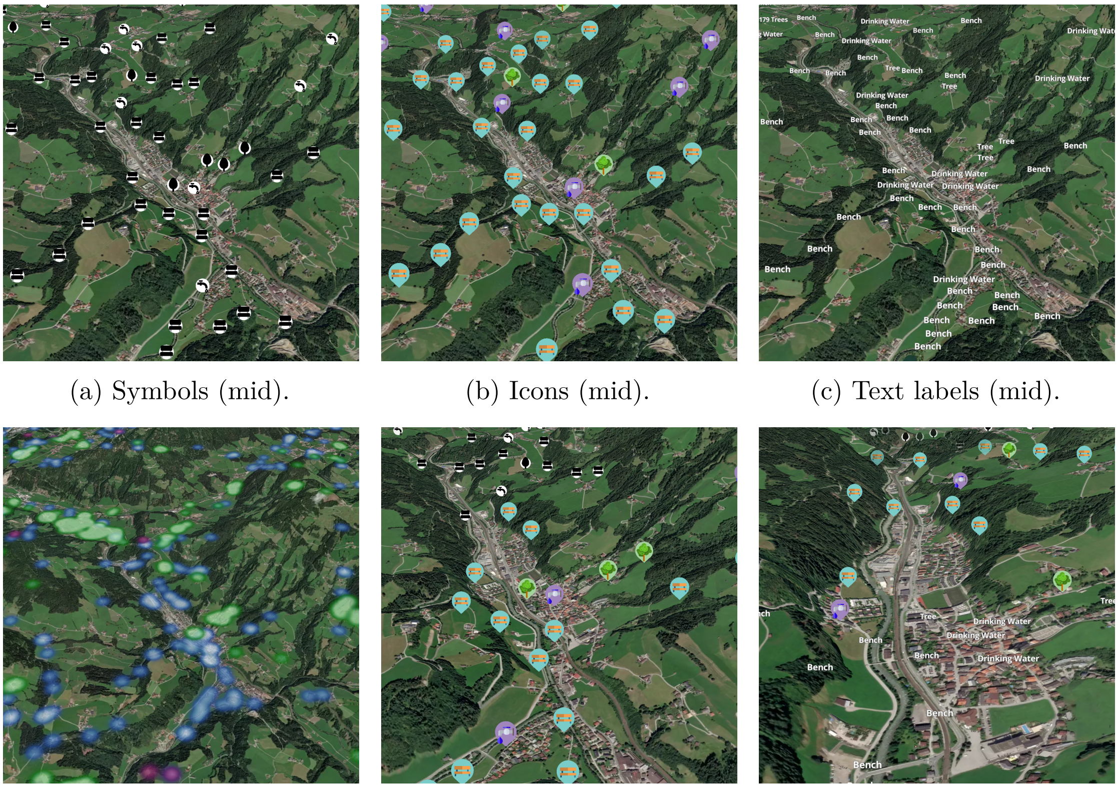

(a) Symbols (mid). (b) Icons (mid). (c) Text labels (mid).

(d) Heatmap (far). (e) Combined (mid). (f) Combined (near).

Figure 5.4: Small rural town.

Findings relative to prior work. *Alpine mountain range* (Fig. 5.3): confirms the benefit of selective realism near ground and refines mid-zoom behavior by showing that symbols remain legible longer on shaded relief than expected from urban contexts, delaying the symbol to icon transition. In contrast, icons and text exhibit insufficient contrast against complex terrain shading, whereas the austere symbol geometry remains easily distinguishable.

Small rural town (Fig. 5.4): confirms effective label-based landmark identification and mid-zoom icon parsing at decision points, highlighting the joint role of visual clarity and semantic density in a compact street network. High PoI density further shows that category-coded icon backgrounds make structural patterns readily perceptible while retaining discrete referents for selection, comparable to the heatmap's macro overview. This underlines the need for density management to mitigate label conflicts and maintain legibility.

Plain (Fig. 5.5): with sparse PoI data and minimal occlusion, icons emerge distinctly at mid zoom and support efficient landmark recognition across open terrain, while labels provide unambiguous identification at near range. Under these conditions, the heatmap contributes limited additional information beyond discrete layers.

(a) Symbols (mid).

(b) Icons (mid).

(c) Text labels (mid).

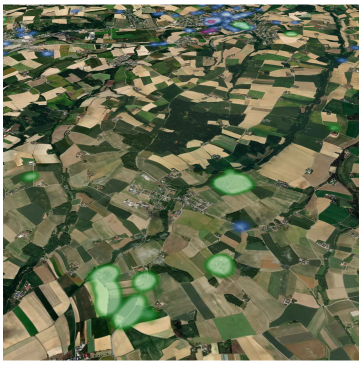

(d) Heatmap (far).

(e) Combined (mid).

(f) Combined (near).

Figure 5.5: Plain.

New hypotheses generated by the inspection:

H1 (Terrain occlusion & LoA thresholds). In environments dominated by terrain-induced occlusion, the transition from symbol to icon visualization occurs at larger distances compared to flat urban areas.

H2 (Category entropy effect). Icons are more effective than text for initial scanning in areas with high category entropy (diverse mix of categories), while text labels perform better for identification tasks in low entropy environments dominated by a single category.

H3 (Backgrounds vs. relief). Category-colored icon backgrounds on icons improve icon parsing at mid-zoom levels but reduce legibility when background luminance contrast is low.

H4 (Heatmap draping effect). Category-specific color enhance heatmap effectiveness at far zoom levels. However, non-draped heatmaps, while suitable for macro overviews, can bias hotspot perception on slopes, highlighting the importance of terrain-aware draping.

Answer. Effective visualization in rural 3D terrain emerges from a balanced blend of realism and abstraction, combined with scale-proportional density control and efficient delivery mechanisms (see Table 5.1). The three case studies support core findings from prior research while generating rural-specific hypotheses (H1–H4) that clarify the circumstances under which each visualization modality is most advantageous. These hypotheses can be empirically tested by systematically varying relief, background luminance, and zoom level, while measuring identification accuracy, time to category decision, and spatial mislocalisation on slopes.

Limitations and Threats to Validity

This evaluation is qualitative and reflective rather than based on a controlled experiment, and several limitations should be considered when interpreting the results.

Case selection and external validity. The analysis covers only three cases (alpine range, small rural town, plain), which restricts the generalizability of findings. Nonetheless, these scenes capture essential variations in relief, semantic density, and occlusion (Figures 5.3–5.5).

Landmark data quality. Landmark selection and category mapping depend heavily on OSM completeness and the consistency of tagging in rural areas. This dependency limits the generalizability of conclusions about landmark visualization layers (see Sections 5.4, 3.2).

LoA selection method. The current combined visualization layer employs distance-only based LoA selection. While relevance- or salience-aware policies have been proposed in prior work, these were not implemented in the prototype (see Sections 2.2, 5.1).

Heatmap terrain draping. Heatmaps were not applied with terrain draping, which led to observed edge misreadings on slopes.

Performance assessment. System performance has not been systematically benchmarked. Metrics such as node-fetch latency, cache hit ratios, and frame times under peak PoI densities remain to be quantified (Section 5.3).

User diversity and validation. While multiple visualization modalities are offered to address diverse user needs, this approach has not been validated through controlled user studies. Several hypotheses point to the importance of stratification by terrain relief and display contrast for future validation.

Implications

Practice. This thesis demonstrates the importance of designing landmark visualizations that balance visual clarity, semantic relevance, and computational efficiency, specifically tailored to rural 3D terrain contexts. Practitioners developing spatial navigation and mapping tools can adopt contrast-aware styling by carefully adjusting icon backgrounds and label outlines to the underlying shaded relief to improve landmark legibility. Terrain-draped heatmaps provide more accurate spatial context and reduce distortion effects typical in slope regions, enhancing macro-scale overview and hotspot detection. Employing a combined distance-based visualization layer minimizes distracting layer switching while retaining the flexibility of discrete layers for in-depth analysis and communication. These design guidelines offer a practical blueprint for improving usability and cognitive support in rural map applications where traditional urban-centric approaches fall short.

Research. This work identifies concrete, testable hypotheses concerning the interaction of terrain relief, visualization styles, and zoom-dependent landmark abstraction. Future research should focus on systematically evaluating these hypotheses through controlled experimental studies that manipulate key factors such as terrain relief, background luminance, and zoom levels while measuring user-focused metrics, including identification accuracy, decision time, and spatial localization errors on slopes. Comparative analysis of distance-only versus salience-aware LoA selection strategies is needed to determine optimal strategies for varying rural contexts. Additionally, performance benchmarking under typical rural bandwidth conditions, especially prefetching and caching behaviors, will support practical deployments. A broader user study incorporating diverse participants and stratifying results by spatial ability and environmental complexity will strengthen the generalizability of these findings. Such empirical efforts will provide essential feedback loops between prototype design and effective, evidence-based landmark visualization in 3D rural navigation.

5.6 Concluding Remarks and Outlook

This chapter has addressed the research question by integrating a layered symbol system with scale-aware delivery and an OSM-based pipeline. The qualitative evaluation across three contrasting rural settings substantiated core findings from prior work and formulated rural-specific hypotheses (H1–H4). The recommended design balances abstraction and realism, applying density control via LoD, supported by principled access and interaction (Table 5.1).

Contributions. This thesis makes several key contributions: an offline retrieval and preprocessing pipeline for OSM-derived PoIs that enables multi-scale overlays in country-scale 3D digital maps; the design and implementation of six PoI representations (symbols, icons, text labels, selected 3D models, combined 2D and a categorized heatmap) integrated into a coherent, scalable visualization stack; and a literature-grounded synthesis interpreting case-based behavior across scales and modalities. The resulting synthesis links observations to prior work and formulates case-specific hypotheses, establishing a concrete agenda for subsequent empirical research.

Future work. Future research should focus on contrast-adaptive icon and label styling, terrain-draped heatmaps, and conditional triggers for 3D representations. Controlled empirical validation should test the hypotheses across varying terrain relief, luminance, and zoom levels, employing metrics such as identification accuracy, decision time, and slope mislocalisation. Additionally, performance benchmarking and comprehensive data quality auditing are vital to facilitating transfer to other regions and applications.

Practical relevance. For rural exploration and wayfinding, the presented visualization stack offers a practical pathway combining macro overview, category search, and near-field confirmation without architectural changes. The combined layer enables direct in situ comparisons, while the discrete layers remain accessible for detailed analysis.

Overview of Generative AI Tools Used

In accordance with the declaration at the beginning of the thesis, generative AI was used solely as an auxiliary aid, and the author's creative and scholarly contribution predominates throughout. Each use case is documented, specifying where and how AI tools were used.

Bibliographic Validation. ChatGPT (GPT-5 Thinking, used on 26.08.2025) was employed to validate every BibTeX entry. It checked the presence and plausibility of core metadata (author(s), title, publisher or venue, year, volume/issue, pages, DOI/ISBN, URL) and flagged missing or malformed fields (e.g., truncated page ranges, inconsistent initials) and contradictions (e.g., DOIs not matching the stated journal or publisher). It did not add sources or rewrite records. Instead, it produced an issues checklist for the author to resolve. Citation styling was handled by LaTeX.

Language Assistance. Grammarly was used during drafting for spell-checking and basic grammar/punctuation alerts. Suggestions were reviewed individually by the author; the tool did not generate or substantially rewrite text. ChatGPT (GPT-5 Thinking) served as a higher-level language checker to flag awkward or non-idiomatic phrasing and unclear sentences. It provided targeted alternatives only; all rewrites were performed and approved by the author, and no substantive content was generated.

List of Figures

List of Tables

Acronyms

2D two-dimensional. 1, 2, 7–10, 14, 19, 21, 29, 33, 35, 49

3D three-dimensional. 1–5, 7–14, 19, 21, 22, 27, 29, 30, 33, 34, 36, 37, 39–42, 44–46, 49

BBox Bounding Box. 18, 22, 23, 26, 30, 39

DEM Digital Elevation Model. 8, 9

LoA Level(s) of Abstraction. 10, 15, 19, 27, 33–35, 40, 41, 44, 45

LoD Level of Detail. 1–4, 10–13, 15–20, 27, 29, 33, 35, 37, 38, 40, 46

OSM OpenStreetMap. 2–5, 13, 15–17, 20–23, 32, 39, 40, 44, 46

PoI Point of Interest. 2, 9, 14–16, 18–23, 26, 28, 31, 32, 43, 45, 46, 49

PoIs Points of Interest. 2, 3, 12, 13, 17, 19, 20, 23–26, 29, 31, 32, 35, 38, 41, 46, 49

UI User Interface. 3, 4, 14, 16, 19, 35, 37

Bibliography

[Al-Bakri, 2015] Al-Bakri, M. (2015). Ten years of OpenStreetMap project: have we addressed data quality appropriately?—review paper. *Journal of Engineering*, 21(4):158–175.

[Battle et al., 2016] Battle, L., Chang, R., and Stonebraker, M. (2016). Dynamic prefetching of data tiles for interactive visualization. In Özcan, F., Koutrika, G., and Madden, S., editors, *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data*, SIGMOD '16, pages 1363–1375, New York, NY, USA. Association for Computing Machinery.

[Carbonell Carrera et al., 2017] Carbonell Carrera, C., Avarvarei, B. V., Chelariu, E. L., Draghia, L., and Avarvarei, S. C. (2017). Map-reading skill development with 3D technologies. *Journal of Geography*, 116(5):197–205.

[Cesium GS, Inc., 2025] Cesium GS, Inc. (2025). Cesiumjs. `https://cesium.com/platform/cesiumjs/`. Accessed: 2025-08-18.

[Cheng et al., 2022] Cheng, B., Wunderlich, A., Gramann, K., Lin, E., and Fabrikant, S. I. (2022). The effect of landmark visualization in mobile maps on brain activity during navigation: A virtual reality study. *Frontiers in Virtual Reality*, 3:981625.

[Cliburn et al., 2007] Cliburn, D. C., Winlock, T., Rilea, S., and Van Donsel, M. (2007). Dynamic landmark placement as a navigation aid in virtual worlds. In Majumder, A., Hodges, L. F., Cohen-Or, D., and Spencer, S. N., editors, *Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology (VRST '07)*, VRST '07, pages 211–214, New York, NY, USA. Association for Computing Machinery.

[Çöltekin et al., 2018] Çöltekin, A., Francelet, R., Richter, K., Thoresen, J., and Fabrikant, S. I. (2018). The effects of visual realism, spatial abilities, and competition on performance in map-based route learning in men. *Cartography and Geographic Information Science*, 45(4):339–353.

[Deakin, 1996] Deakin, A. K. (1996). Landmarks as navigational aids on street maps. *Cartography and Geographic Information Systems*, 23(1):21–36.

[Elias and Brenner, 2005] Elias, B. and Brenner, C. (2005). Automatic generation and application of landmarks in navigation data sets. In Fisher, P. F., editor, *Developments in Spatial Data Handling: 11th International Symposium on Spatial Data Handling*, pages 469–480, Berlin, Heidelberg. Springer, Berlin, Heidelberg.

[Elias and Paelke, 2008] Elias, B. and Paelke, V. (2008). *User-Centered Design of Landmark Visualizations*, pages 33–56. Lecture Notes in Geoinformation and Cartography. Springer, Berlin, Heidelberg.

[Franke and Schweikart, 2017] Franke, C. and Schweikart, J. (2017). Mental representation of landmarks on maps: Investigating cartographic visualization methods with eye tracking technology. *Spatial Cognition & Computation*, 17(1-2):20–38.

[Gaede and Günther, 1998] Gaede, V. and Günther, O. (1998). Multidimensional access methods. *ACM Computing Surveys (CSUR)*, 30(2):170–231.

[jscastro76, 2025] jscastro76 (2025). Threebox — mapbox gl js + three.js integration. [Online; accessed 20-September-2025; version 2.2.7].

[Kapaj et al., 2024] Kapaj, A., Hilton, C., Lanini-Maggi, S., and Fabrikant, S. I. (2024). The influence of landmark visualization style on task performance, visual attention, and spatial learning in a real-world navigation task. *Spatial Cognition & Computation*, 24(4):227–267.

[Kapaj et al., 2022] Kapaj, A., Lin, E., and Lanini-Maggi, S. (2022). The effect of abstract vs. realistic 3D visualization on landmark and route knowledge acquisition. In Ishikawa, T., Fabrikant, S. I., and Winter, S., editors, *15th International Conference on Spatial Information Theory (COSIT 2022)*, volume 240 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:8, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[Keil et al., 2020] Keil, J., Edler, D., Kuchinke, L., and Dickmann, F. (2020). Effects of visual map complexity on the attentional processing of landmarks. *PLOS ONE*, 15(3):e0229575.

[Lee et al., 2001] Lee, Y. C., Kwong, A., Cheng, S., and Mack, A. (2001). Multi-media map for visual navigation. *Journal of Geospatial Engineering*, 3(2):87–96.

[Lindstrom et al., 1996] Lindstrom, P., Koller, D., Ribarsky, W., Hodges, L. F., Faust, N., and Turner, G. A. (1996). Real-time, continuous level of detail rendering of height fields. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 109–118. Association for Computing Machinery.

[Mapbox, 2025] Mapbox (2025). Mapbox GL JS. `https://www.mapbox.com/mapbox-gljs`. Accessed: 2025-08-18.

[MapLibre, 2025] MapLibre (2025). Maplibre gl js. `https://maplibre.org/maplibre-gl-js/`. Accessed: 2025-09-10.

[Maps-for-Free, 2025] Maps-for-Free (2025). Digital elevation model (dem) png. Screenshot downloaded from maps-for-free.com. Online; accessed 18-August-2025, Released under Creative Commons CC0—no attribution required.

[Microsoft, 2025] Microsoft (2025). Typescript — typed javascript. [Online; accessed 20-September-2025; version 5.7.2].

[Nurminen and Sirvio, 2021] Nurminen, A. O. I. and Sirvio, K. M. (2021). Bus stop spotting: a field experiment comparing 2D maps, augmented reality and 3D maps. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction*, MobileHCI '21, New York, NY, USA. Association for Computing Machinery.

[OpenJS Foundation, 2025] OpenJS Foundation (2025). deck.gl. `https://deck.gl/`. Accessed: 2025-08-18.

[OpenStreetMap Foundation, 2025] OpenStreetMap Foundation (2025). Openstreetmap. `https://www.openstreetmap.org`. Online; Accessed: 2025-08-18.

[OpenStreetMap Wiki, 2024] OpenStreetMap Wiki (2024). Bounding box — openstreetmap wiki. [Online; accessed 20-September-2025].

[OpenStreetMap Wiki, 2025] OpenStreetMap Wiki (2025). Overpass api — openstreetmap wiki. [Online; accessed 18-August-2025].

[Pasewaldt et al., 2012] Pasewaldt, S., Semmo, A., Trapp, M., and Döllner, J. (2012). Towards comprehensible digital 3d maps. In Jobst, M., editor, *SOMAP 2012: Proceedings of the Symposium on Service-Oriented Mapping, 22–23 November 2012, Vienna*, pages 261–276, Vienna. JOBSTMedia.

[Semmo et al., 2012] Semmo, A., Trapp, M., Kyprianidis, J., and Döllner, J. (2012). Interactive visualization of generalized virtual 3d city models using level-of-abstraction transitions. *Computer Graphics Forum*, 31:885–894.

[Smallman and Cook, 2011] Smallman, H. S. and Cook, M. B. (2011). Naïve realism: Folk fallacies in the design and use of visual displays. *Topics in Cognitive Science*, 3(3):579–608.

[Three.js, 2025] Three.js (2025). three.js — javascript 3d library. [Online; accessed 20-September-2025; version 0.170.0].

[Webpack, 2025] Webpack (2025). webpack — a bundler for javascript. [Online; accessed 20-September-2025].

[Yesiltepe et al., 2021] Yesiltepe, D., Conroy Dalton, R., and Özbil Torun, A. (2021). Landmarks in wayfinding: a review of the existing literature. *Cognitive Processing*, 22(3):369–410.

[Zhu et al., 2022] Zhu, L., Shen, J., Zhou, J., Stachoň, Z., Hong, S., and Wang, X. (2022). Personalized landmark adaptive visualization method for pedestrian navigation maps: Considering user familiarity. *Transactions in GIS*, 26(2):669–690.