



**PAISII HILENDARSKI
UNIVERSITY OF PLOVDIV**



**FACULTY OF
MATHEMATICS
AND INFORMATICS**

RECTORATE:

24 Tsar Asen Str., Plovdiv 4000, Bulgaria

phone: +359 32 261 261

email: rector@uni-plovdiv.bg

uni-plovdiv.bg

DEANERY:

236 Bulgaria Blvd., Plovdiv 4027, Bulgaria

phone: +359 32 261 800

email: dean-fmi@uni-plovdiv.bg

fmi-plovdiv.org

Department of Computer Systems

Research on the Development of an Ambient-Oriented Route Generator

ABSTRACT

for the award of the educational and scientific degree “Doctor”
in field 4 of higher education: Natural Sciences, Mathematics and Informatics,
professional direction: 4.6. Informatics and Computer Science,
doctoral programme: Informatics

PhD Candidate: **Georgi Kostadinov Kostadinov**

Scientific Supervisor: **Prof. Asya Stoyanova-Doycheva, PhD**

Plovdiv
2026

The abstract has been prepared on the basis of the dissertation “Research on the Development of an Ambient-Oriented Route Generator”, developed at the Department of Computer Systems within the Faculty of Mathematics and Informatics of the *Paisii Hilendarski* University of Plovdiv, city of Plovdiv.

The dissertation comprises 157 pages, 10 figures, 21 tables and 6 appendices. The bibliography includes 83 sources. The list of the author’s publications on the topic consists of 4 titles.

The data on the meeting, the composition of the scientific jury, and the date and place of the defence are to be filled in after the procedure has been approved.

Scientific jury:

1.
2.
3.
4.
5.

Author: Georgi Kostadinov Kostadinov

Contents

Introduction.....	4
Objectives of the dissertation.....	5
Structure of the dissertation	5
Chapter 1. Current state of the art.....	6
1.1 Recommender systems in tourism	6
1.2 Interactive tourist guides.....	6
1.3 Ontologies, context and personalization in tourism systems	6
1.4 Accessibility, explainability and uncertainty in intelligent tourism	7
1.5 Comparative analysis of existing solutions.....	7
1.6 Conclusions of Chapter 1	8
Chapter 2. Architecture of an ambient-oriented intelligent tourist guide	9
2.1 Main functionalities of the intelligent tourist guide.....	9
2.2 Virtual physical space	9
2.3 Architecture of the tourist guide	10
2.4 Conclusions of Chapter 2	13
Chapter 3. Implementation of a prototype of an ambient-oriented intelligent tourist guide	14
3.1 Technical implementation.....	14
3.2 Operating logic of the tourist guide	16
3.3 Limitations of the prototype.....	18
3.4 Toolkit for expert validation	18
3.5 Analysis of the influence of individual ambient factors	18
3.6 Extensions of the data and knowledge model.....	21
3.7 Ethical and applied aspects of the proposed solution	21
3.8 Conclusions of Chapter 3	21
Chapter 4. Experimental validation of the proposed model	22
4.1 Experimental setup.....	22
4.2 Quantitative results from the comparison of strategies.....	22
4.3 Ablation analysis of the ambient factors.....	23
4.4 Scenario validation and qualitative analysis	23
4.5 Discussion of results and limitations of the validation	23
4.6 Validation with real data from the subject domain	23
4.7 Conclusions of Chapter 4.....	24
Conclusion	26
Results of the dissertation	26
Directions for future research	27
Publications on the topic	28
References.....	30

Introduction

When a tourist decides to visit an unfamiliar city or region, they face a practical problem: how to choose which cultural and historical sites to visit, in what order and under what conditions, so that the visit is both feasible and satisfying. This problem remains unsolved despite the growing volume of available information about cultural heritage — descriptions, categories, locations and so on — which by itself is not sufficient for making a decision. Today the user no longer expects merely access to a catalogue of tourist sites, but an intelligent environment that helps them make an informed choice, assess the key constraints of the specific situation and obtain a practically feasible route¹. This dynamic is particularly pronounced in mobile tourist applications [1], [2], in which the decision of which next site to choose is strictly contextual and is taken on the spot, under the influence of multiple variables: time constraints, spatial proximity, physical accessibility and individual user preferences [3], [4], [5], [6], [7], [8].

Classical electronic tourist guides traditionally provide list-based or map-based representations of attractions, a brief description, images and sometimes basic category filters. Although such systems have high reference value, they leave the user with the difficult task of combining on their own the thematic interest, the current location, the available time, their personal constraints and the changing conditions of the environment. Their informational value remains high, but they exhibit a deficit of contextual intelligence. The present work realizes the notion of “ambient characteristics” — dynamic properties of the sites and the environment around them (location, accessibility, weather conditions, exposure and time availability) — which determine the practical suitability of a given visit at a particular moment. From this arises the need for a new generation of tourist guides capable of integrating semantic knowledge representation, dynamic context awareness and an algorithm for personalized decision making [9], [10], [11], [12], [13]. The scientific and applied problem of the present work is thus outlined: how to design a tourist guide that simultaneously integrates semantic representation of cultural and historical sites, dynamic context awareness and personalized decision making. The ultimate aim is the generation of practically feasible and explainable itineraries under conditions of variable external context and diverse user profiles.

In the scientific literature the problem of tourist recommendation is examined along several intersecting lines. Recommender systems use content-based, collaborative and hybrid models for selecting sites according to the user’s profile [11], [14], [15], [1]. Mobile guides and context-aware applications use location, time and behavioural signals in order to provide more relevant content [10], [8], [16]. Ontological and semantic modelling of the tourism domain enables a formal description of cultural sites, their categories, properties and relationships [17], [18], [19]. In addition to these three lines, interest is also growing in accessible tourism, the explainability of recommendations and the handling of incomplete data, which raise the question of the applicability of and trust in intelligent tourism systems [20], [21], [22], [23], [24].

The present dissertation is positioned precisely at the intersection of these lines. It addresses the problem of an ambient-oriented intelligent tourist guide, in which tourist sites are not realized as purely static.

The thesis of the work is that the intelligent tourist guide should be modelled as a context-sensitive, semantically informed and explainable recommendation environment. In such an environment, the initial selection of candidate sites can be performed through several equivalent paths: spatial discovery of nearby sites, search by ontological concepts, navigation to similar sites on

¹ In the present work, a “route” is understood as a prioritized list of candidate sites for visiting, ordered by their degree of contextual suitability, rather than as a geographical path of movement.

the basis of semantic closeness, or explicit user choice. All these paths feed candidates into a shared ambient re-ranking component that integrates both the external context and the explicit user profile.

Objectives of the dissertation

The main objective of the dissertation is to design, implement and substantiate a two-stage model of an ambient-oriented intelligent tourist guide, in which a semantic retrieval of candidates based on ontological knowledge is followed by an ambient re-ranking that takes into account the user profile and the dynamic external context, with the aim of generating practically feasible and explainable routes (in the sense defined above) to cultural and historical sites.

To achieve this objective, three main sub-objectives are formulated, elaborated in specific research tasks:

- Objective 1. To develop a theoretical and architectural framework for an ambient-oriented intelligent tourist guide, in which the semantic description of tourist sites is combined with a multidimensional model of the ambient factors.
 - Task 1.1. To systematize and analyze the current state of scientific work in the field of tourist recommender systems, context-aware tourist guides, ontological models and accessible tourism, in order to identify the research niche in which the work is positioned.
 - Task 1.2. To define a formal model of a tourist site as an ontological instance extended with ambient characteristics, and to build a multi-layered architecture compatible with the principles of the virtual physical space, which clearly separates the stage of semantic retrieval from the stage of ambient re-ranking and enables the integration of knowledge, user profile, services and external context.
- Objective 2. To implement a prototype system that materializes the two-stage model in a working software project with support for semantic search, ambient re-ranking, explainability and uncertainty management.
 - Task 2.1. To implement the main modules of the prototype — search by ontological concepts, retrieval of similar and nearby sites, management of user profiles and ratings, ambient re-ranking algorithm — and to integrate a dynamic meteorological context so that the system reacts to environmental conditions in real time.
 - Task 2.2. To formalize accessibility (motor, visual, auditory) as a real constraint in ambient re-ranking, rather than a descriptive attribute.
 - Task 2.3. To propose a cautious mechanism for controlled completion of missing or incomplete ambient characteristics through a language model and the possibility of expert verification, which raises the coverage of applicable routes without sacrificing reliability.
 - Task 2.4. To embed recommendation explainability and uncertainty management as inalienable qualities of the guide.
- Objective 3. To formulate and apply a methodology for experimental validation of the proposed model, including comparison with baseline recommendation strategies, ablation analysis of the individual ambient factors, scenario-based validation by user profile types, and verification on both a synthetic and a real corpus from the subject domain.

Structure of the dissertation

The dissertation moves from theoretical foundation (Chapter 1 — current state of the field and the research niche) through the architectural design (Chapter 2) and the engineering implementation of the prototype (Chapter 3) to the experimental validation on synthetic and real corpora (Chapter 4); the Introduction derives the goals and tasks, while the Conclusion summarizes the results, limitations, future directions and publications.

Chapter 1. Current state of the art

1.1 Recommender systems in tourism

Recommender systems encompass content-based [15], collaborative [11] and hybrid [14] approaches. In tourism, however, this general categorization proves insufficient, because the sites are simultaneously spatial, temporal and semantic units, and their applicability depends on distance, working hours and momentary constraints. Ricci [25] outlines the specifics of recommender systems for travel as a separate class, Gavalas et al. [1] later systematize contemporary mobile tourist recommender systems, while earlier works on route planning [9], [26], [27] introduce the formal constraints of time windows and profits.

Content-based models are useful in the “cold start” situation, when sites have rich semantic descriptions [18], [15], [19]. For cultural sites, however, metadata are rarely complete and homogeneous, which makes these models fragile with respect to dynamic context and real applicability.

Scenario: a tourist with motor impairments looks for sites on a rainy Sunday afternoon. A system without accessibility knowledge may recommend an open-air archaeological reserve — thematically relevant but physically infeasible; a system with accessibility data but no weather context will offer an accessible open-air site without accounting for the rain. Only when accessibility, weather, working hours and proximity are evaluated simultaneously can the system eliminate inapplicable candidates and prioritize the truly feasible ones.

1.2 Interactive tourist guides

An interactive tourist guide combines mobile access, localization and adapted content, and early systems such as Cyberguide [8] laid precisely this foundation. For solving the actual tourist choice, however, this is not sufficient, because location must be supplemented with a richer context and a mechanism for assessing applicability [10], [16], [28].

Interactive guides combine search, navigation and recommendation around the user’s movement in space [3], [4], [5] and are increasingly part of broader intelligent architectures with services, knowledge bases and external sources [29], [30]. This functional diversity in itself does not guarantee intelligent behaviour, unless the system formally models the knowledge and the dynamic context of the choice.

From the standpoint of the dissertation problem, the most essential point is that interactive guides show how a tourist system can be placed in a real context of use, but they do not always provide a sufficiently strong model for explainable and context-aware recommendation. The ambient approach extends the idea of the interactive guide by introducing formalized characteristics of the environment and of the suitability of the site for a visit.

1.3 Ontologies, context and personalization in tourism systems

An ontology provides a formal description of tourist sites and the relationships between them, so that the knowledge is optimized for search and reusable [17], [18], [19]. The semantic description alone, however, does not resolve whether the site is applicable at the current moment for the specific tourist.

Ontological and taxonomic models allow a single site to be described along several semantic axes, while the context links this stable knowledge with the specific use [17], [30]. The critical point is that without a formally represented context the ontology remains good for search but insufficient for a dynamic decision in a real tourist situation.

1.4 Accessibility, explainability and uncertainty in intelligent tourism

The term “smart tourism” is used in the present work in the sense given by Gretzel et al. [31], as a tourist environment in which information and communication technologies, together with data, transform services into adaptive, context-sensitive and personalized ones. Within this framework, accessibility, explainability and the management of uncertainty emerge as fundamental, but insufficiently jointly developed, requirements. Accessibility cannot be modelled as a single binary attribute; it requires at least a distinction between different types of constraints and forms of disadvantage [20], [23], [24]. Without such a distinction, the tourist system risks recommending sites that are formally close but practically infeasible [32].

The literature on routing for users with motor impairments shows that spatial proximity by itself does not guarantee the practical applicability of the route [20], [23]. The route must not merely be short, but physically feasible and consistent with the user’s constraints. Analogously, in the tourism domain a site must not be recommended solely because it is nearby or popular if the conditions for visiting it are absent. This makes accessibility a structural rather than a peripheral factor.

Combining accessibility, explainability and uncertainty management raises the bar for an intelligent tourist guide: the system must not only find good candidates but do so transparently, fairly and robustly to incomplete data — a distinction that sets the ambient-oriented approach apart from simpler recommendation strategies.

1.5 Comparative analysis of existing solutions

On the basis of the approaches considered so far, the present section proposes a systematized comparative analysis of leading existing systems along the five axes by which the result of the dissertation is delineated: semantic description of the site, multidimensional accessibility, dynamic weather context, controlled completion via AI and explainability of the recommendation.

The choice of these five axes is not arbitrary. The semantic description has been identified as a necessary condition for reusability and thematic search in numerous independent surveys [17], [18], [19]. Accessibility is recognized as an autonomous research direction in the tourism domain [23], [24], [32]. The dynamic weather context is identified as a significant factor for tourist behaviour [33], [34], [35]. Controlled data completion addresses the documented problem of incompleteness of cultural and historical metadata [36], [37], [38]. Explainability is established as a separate taxonomy within recommender systems [21], [22], [39]. The following table additionally includes an axis for global route optimization, on which existing systems surpass the present prototype, in order to ensure a balanced comparison.

Table 1. Summary positioning of the proposed model with respect to existing solutions.

Criterion	Typical solutions	Proposed model
Semantics	Categories, tags or ontologies	Ontological concepts and instances
Accessibility	Descriptive attribute	Hard constraint in case of incompatibility
Weather context	Rarely dynamic	Exposure with current weather
Explainability	Often absent	Factor explanations and reason for elimination
Missing data	Ignored	AI inference with expert verification

For a more precise positioning of the proposed model with respect to related classes of systems, it is useful to consider several conceptually similar types of solutions described in the literature and in practice. Such a comparison shows not only what the present prototype does, but also how it differs from other types of tourism and recommender systems.

1.6 Conclusions of Chapter 1

The analysis of the current state of the field allows several summarizing conclusions to be formulated.

Tourist recommender systems are multi-criteria and context-dependent; existing guides use location and mobility but do not sufficiently integrate a rich semantic model, multidimensional accessibility and handling of incomplete data, while ontological models are an appropriate foundation that must be augmented with a contextual and user layer. Accessibility and explainability are structural requirements, not peripheral properties, and missing data and uncertainty require cautious mechanisms for enrichment and validation, including careful use of AI services.

The analysis outlines five concrete research opportunities: (1) semantic description of sites through ontological concepts and instances as a basis for thematic search and reusability; (2) multidimensional accessibility, treated as a hard constraint; (3) dynamic weather context modelled through the ambient characteristic of exposure; (4) controlled completion of incomplete data through AI services with protective strategies; (5) explainability of ranking through traceable factor contributions and an explicit reason for elimination. These five opportunities define the scope within which the architecture of Chapter 2 and the implementation of Chapter 3 localize the specific research contributions of the dissertation.

Chapter 2. Architecture of an ambient-oriented intelligent tourist guide

2.1 Main functionalities of the intelligent tourist guide

The architecture addresses the five research opportunities from §1.6: semantic description — through the knowledge layer; multidimensional accessibility — through the ambient characteristics; dynamic weather context — through the integration layer; controlled AI completion — through an isolated module; explainability — through traceable orchestration. The system covers the full tourist-support cycle (registration, semantic search, ambient re-ranking, feedback) as a sequence of interrelated operations, not isolated screens or API calls.

The first group of functionalities covers user identification and profiling: registration, storing basic coordinates, available time window, preferences and constraints. The profile is the main input structure for the model — from it, the weights of ambient factors, spatial coordinates and constraints that participate in ranking are derived. The data are obtained through a questionnaire on the client side with the approach of explicit profiling in an intelligent personal tourist guide.

The second group of functionalities encompasses the semantic search and the access to the ontological knowledge base. The user must be able to search concepts by type and keyword, retrieve instances by a chosen concept, view the details of a particular site and find similar sites by shared semantic characteristics. This functional group plays the role of a thematic filter and prepares the initial set of candidates for subsequent ambient evaluation.

The functionalities defined in this way show that the proposed tourist guide is an integrated intelligent environment, in which knowledge, profile, context and algorithmic logic interact within a common architectural model.

2.2 Virtual physical space

The concept of a virtual physical space fits within the broader research field of the Internet of Things [40] and of context-aware computing [41], and at the same time offers a concrete architectural form for combining the physical world, the digital representation and the social interaction in a single applied environment [29], [30]. In the context of the intelligent tourist guide, this means that the system must be able to connect the real location of the tourist, the real properties of the cultural sites and the digital knowledge about them in a single operational cycle. Without such a framework, the tourist guide remains either a purely reference system or a purely geolocation tool.

The first explicit formulation of ViPS as a reference architecture for IoT applications is given in [42], where the space is described as an architecture capable of adapting to different domains and considered in close connection with cyber-physical and cyber-physical-social spaces. This line is developed in greater detail in [43], where ViPS is presented by the authors as the first general reference architecture of an IoT ecosystem, and in [44], where the transition from DeLC and Virtual Education Space to Virtual Physical Space is systematized as a more general model for integrating the virtual, the physical and the social world.

Architecturally, ViPS is not reduced to an abstract idea of connectivity, but to a concrete composition of formalisms and active components. The summary description in [29] and the earlier publications [43], [45] and [46] show that the core of the environment includes the Event Engine/ENet for the representation and dissemination of events, TNet for the temporal aspects, AmbiNet for spatial modelling, personal and operational assistants for rational behaviour and a guard layer for the interface with the IoT nodes. Of particular importance is the work [46], because it specifies how events can be categorized, serialized, packaged and routed through different message brokers without binding the business logic to a particular transport technology.

In the present architecture the virtual-physical space is used as a design principle: the knowledge base contains the ontological concepts and instances (virtual model), the ambient

characteristics link knowledge with the physical feasibility of the visit, and the user profile anchors behaviour to the real tourist. The route is a dynamic product of the interaction between knowledge about the site, context about the user and dynamic data about the environment.

2.3 Architecture of the tourist guide

The proposed tourist guide follows a multi-layered architectural model, in which the individual components are organized around clearly delineated responsibilities in the spirit of Clean Architecture [47] and of domain-driven design [48]. At the most general level, a presentation layer, an application layer, a knowledge layer and an integration layer can be distinguished. This decomposition is chosen not only on engineering grounds [49], but also for the scientific need that the individual research ideas be localized in concrete architectural units. The mapping with respect to the ViPS framework of Section 2.2 is as follows: the knowledge layer realizes the ontology-ambient core, corresponding to the AmbiNet component [45]; the application layer plays the role of an operational assistant, analogous to the assistants of the tourist guide in [3] and [29]; the integration layer corresponds to the guard layer towards external services and events [43], [46]; and the presentation layer provides a technologically neutral REST interface to this environment [50]. Semantic search, ambient ranking, user profiling, weather integration and AI inference are different tasks and must be encapsulated so that the interaction between them is controllable. In the implementation, these logical layers are organized as vertical slices by domain context — each context (ambient, ontology, users) contains its own REST resources, services, models and repositories, instead of being grouped in horizontal packages by technical criterion.

In the knowledge layer, ontological concepts and instances are mapped to separate JPA entity classes and relational tables: `OntologyConcept` is stored in `ontology_concept`, and `OntologyInstance` — in `ontology_instance`. Both types have an external identifier, name, description and a hierarchical relationship `parent_id`; the many-to-many relationship between concepts and instances is realized through the intermediate table `ontology_instance_ontology_concept`. In this way, one concept can group together multiple tourist sites, and one instance can belong to more than one semantic category. The ambient characteristics are linked to the specific instance through `ontology_instance_id` and are stored in a base table `ambient_characteristic` with specialized tables for `location`, `accessibility`, `exposure` and `working_hours`. [47] [48] [49] [45] [3] [29] [43] [46] [50]

The term “ontological model” is used in the present work in an engineering and applied sense. The model contains hierarchically organized concepts, instances and relationships between them, but does not claim to be a complete formal OWL/RDF ontology. From a strictly logical point of view, part of the structure can be regarded as a domain taxonomy, enriched with links to instances and ambient characteristics. This does not diminish the added value of the model, because it is precisely this structured semantics that enables search by concepts, grouping of sites, retrieval of similarity and subsequent ambient re-ranking. [47] [48] [49] [45] [3] [29] [43] [46] [50]

The mapping of the proposed architecture to ViPS is conceptual and not technologically literal. The original ViPS architecture envisages asynchronous dissemination of events through the Event Engine and temporal modelling through TNet, implemented over a message broker infrastructure [46]. In the present system these functions are adapted to a synchronous request–response model: the orchestration of contextual changes (weather data, AI inference, user profile) is carried out within a single HTTP request, rather than through a publish–subscribe mechanism. This choice is dictated by the nature of the application — the tourist guide serves individual requests with a specific user context, rather than a stream of asynchronous sensor events. Accordingly, ViPS is used as a design principle for the separation of responsibilities, rather than as a concrete technological specification.

The integration layer connects the system with external services: a weather source for current conditions and an AI service for inferring missing characteristics. From an architectural point of view, this is a layer with lower reliability compared with the local database, which is why access to it is isolated through dedicated services, fallback behaviour (fallback) and protective mechanisms.

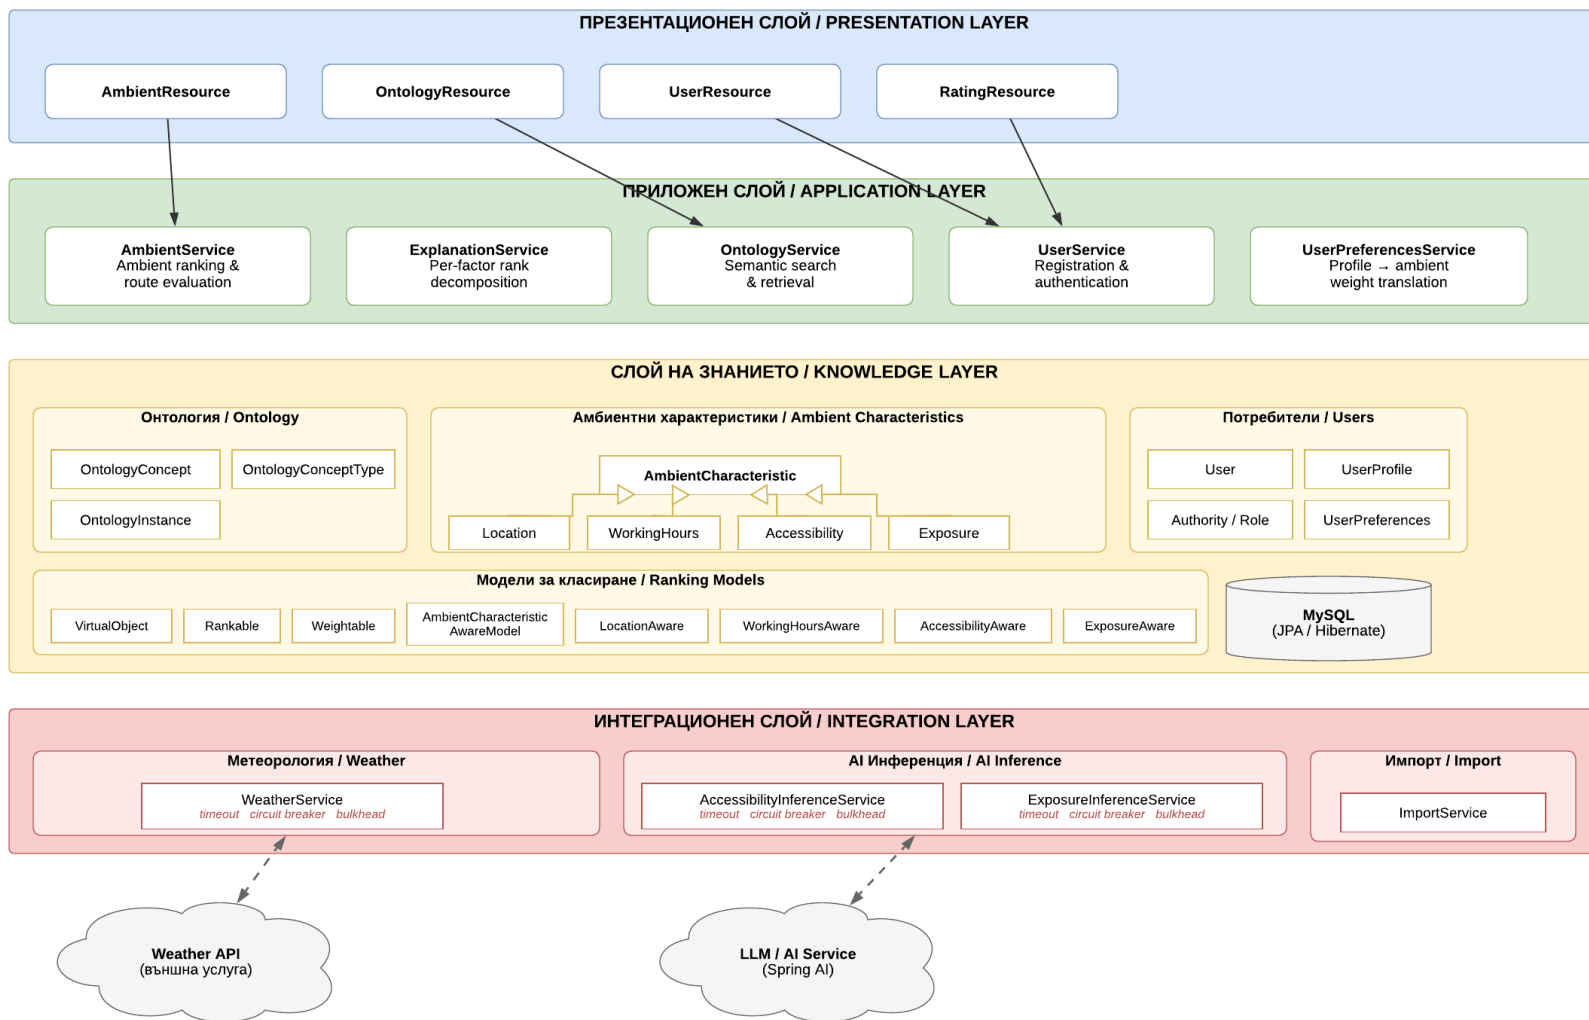


Figure 1. Block diagram of the architectural layers and main subsystems of the ambient-oriented tourist guide.

2.4 Conclusions of Chapter 2

The proposed architecture shows that the ambient-oriented tourist guide can be implemented as a multi-layered intelligent environment, rather than as a monolithic algorithm. Semantic search, ambient evaluation, the user profile and the external context are separated into autonomous architectural units, which facilitates both the extension of the system and its analytical description.

The architecture is consistent with the principles of the virtual physical space, since it links the digital description of the sites with the real conditions for visiting them, and rests on established architectural principles — Clean Architecture [47], domain-driven design [48] and stabilizing patterns (circuit breaker, bulkhead, timeout) for external integration [51]. In this way, the results of the dissertation are localized in concrete architectural units, rather than remaining at a conceptual level.

The first essential result is the ontology-ambient core in the knowledge layer. Through the common base table with specializations (class table inheritance [52]) the ambient characteristics can be processed polymorphically, while at the same time each characteristic retains its own evaluation model. This provides a direct architectural realization of the requirement for extensibility and resolves the problem of multidimensional accessibility formulated in Section 1.6.

Finally, the modular decomposition provides a good basis for scientific defence, because it allows the results to be clearly localized: ontological model, ambient layer, contextual integration and mechanism for handling incomplete data. The same decomposition creates the natural transition to Chapter 3, in which each of these architectural units is examined at the level of data models, algorithms and concrete rules for ambient ranking.

Table 2. Architectural layers and concrete implementation.

Layer	Role	Implementation
Presentation	Access to functions	REST API for the client
Application	Orchestration	AmbientService, OntologyService, ExplanationService
Knowledge base	Semantic model	OntologyConcept and OntologyInstance
Ambient	Physical and situational constraints	Location, Accessibility, Exposure, WorkingHours
Integration	External services	WeatherService and AI inference

Chapter 3. Implementation of a prototype of an ambient-oriented intelligent tourist guide

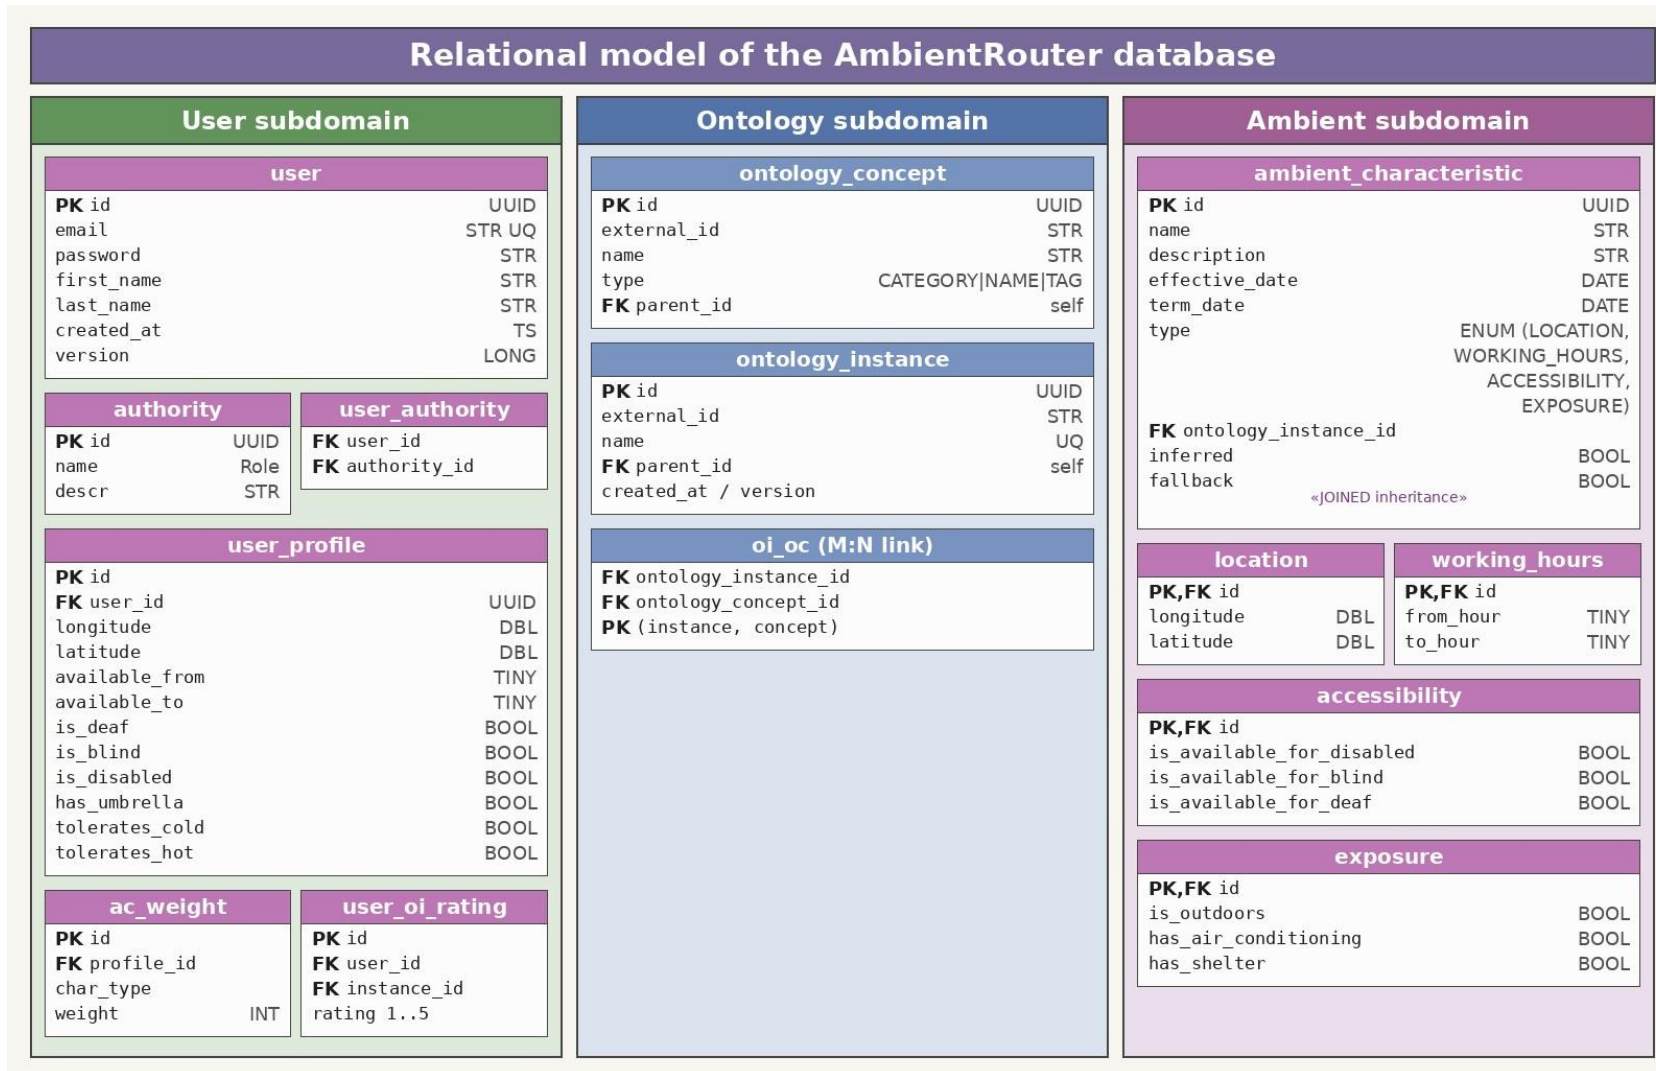
3.1 Technical implementation

The technical implementation is a server-side Java/Spring core exposing a REST interface (Java 21, Spring Boot 3.5, MySQL 8.0). The Spring ecosystem provides a clean separation between the web layer, domain services, data repositories and integrations, matching the architectural setup of the previous chapter and facilitating the application of the scientific ideas.

The prototype is organized by domain — `domain/ontology`, `domain/ambient`, `domain/user` — plus infrastructure modules for weather integration, AI inference and import. The ontology model, ambient characteristics and user profile are thus actually implemented as separate interacting modules of the system.

Built into the model of the exchanged data is the principle of traceability: the system distinguishes directly available values from inferred ones and allows the origin of the individual factors to be shown (directly entered, inferred or the product of the fallback mechanism).

The data model is a relational schema with three sub-domains: an ontology layer (`ontology_instance`, `ontology_concept` and a link table); an ambient layer (common `ambient_characteristic` table with specializations for `accessibility`, `exposure`, `location`, `working_hours` — polymorphic treatment with distinct scoring models); and a user domain (identity plus algorithmic parameters — coordinates, time window, constraints and factor weights).



Ambient subdomain

ambient_characteristic

PK id

UUID

name

STR

description

STR

effective_date

DATE

term_date

DATE

type

ENUM (LOCATION, WORKING_HOURS, ACCESSIBILITY, EXPOSURE)

FK ontology_instance_id

inferred

BOOL

fallback

BOOL

«JOINED inheritance»

location

PK,FK id

longitude

DBL

latitude

DBL

working_hours

PK,FK id

from_hour

TINY

to_hour

TINY

accessibility

PK,FK id

is_available_for_disabled

BOOL

is_available_for_blind

BOOL

is_available_for_deaf

BOOL

exposure

PK,FK id

is_outdoors

BOOL

has_air_conditioning

BOOL

has_shelter

BOOL

Figure 2. ER diagram of the relational model of the AmbientRouter database.

The user context is collected at registration through a questionnaire, implemented in the client application. The explicit profiling approach is the one described by Kristian Milev [53] in the context of an intelligent personal tourist guide. The collected data — coordinates, time window, constraints and weights of the ambient factors — feed directly into the ranking algorithm, which ensures personalization from the very first use.

AI inference is integrated with the deterministic ambient model as structural enrichment under managed uncertainty, not as a replacement of the deterministic logic. Three specialized services are implemented — `AccessibilityInferenceService`, `ExposureInferenceService` and `WorkingHoursInferenceService` — each with a structured prompt, low temperature and a clearly defined output format, which transfers AI inference from free text to a constrained structural channel.

For accessibility the model extracts three separate values (motor, visual, auditory); for exposure — `isOutdoors`, `hasAirConditioning` and `hasShelter`; for working hours — start and end hour. Unresolved fields receive a neutral rank (1.0) uniformly across all three characteristics — this leads to neutral treatment, rather than to pessimistic fallback values.

The language-model context is constructed from the description of the ambient and the available ontological concepts of the associated site: the ontology provides structured domain context, while the language model performs probabilistic completion of missing attributes. The AI output is treated as a conditionally inferred value, distinguishable from observed data and subject to subsequent human verification.

3.2 Operating logic of the tourist guide

The operating logic of the guide follows a two-stage model. At the first stage, an initial set of candidate sites is formed through semantic search by ontological concepts, retrieval of nearby sites around the user, or selection of similar sites with respect to an instance already being viewed. At the second stage, this set is interpreted with respect to the current profile and environment, so that the system answers not only which sites are thematically meaningful, but which are applicable right now.

Intermediate results from the development of this two-stage approach have been published in [7], where an algorithm is presented for generating and visualizing the route of an intelligent tourist guide, and in [6], where the model is extended with ambient characteristics, personalization and ontological concepts. The present work builds upon these results by integrating them in an overall architecture and validation methodology.

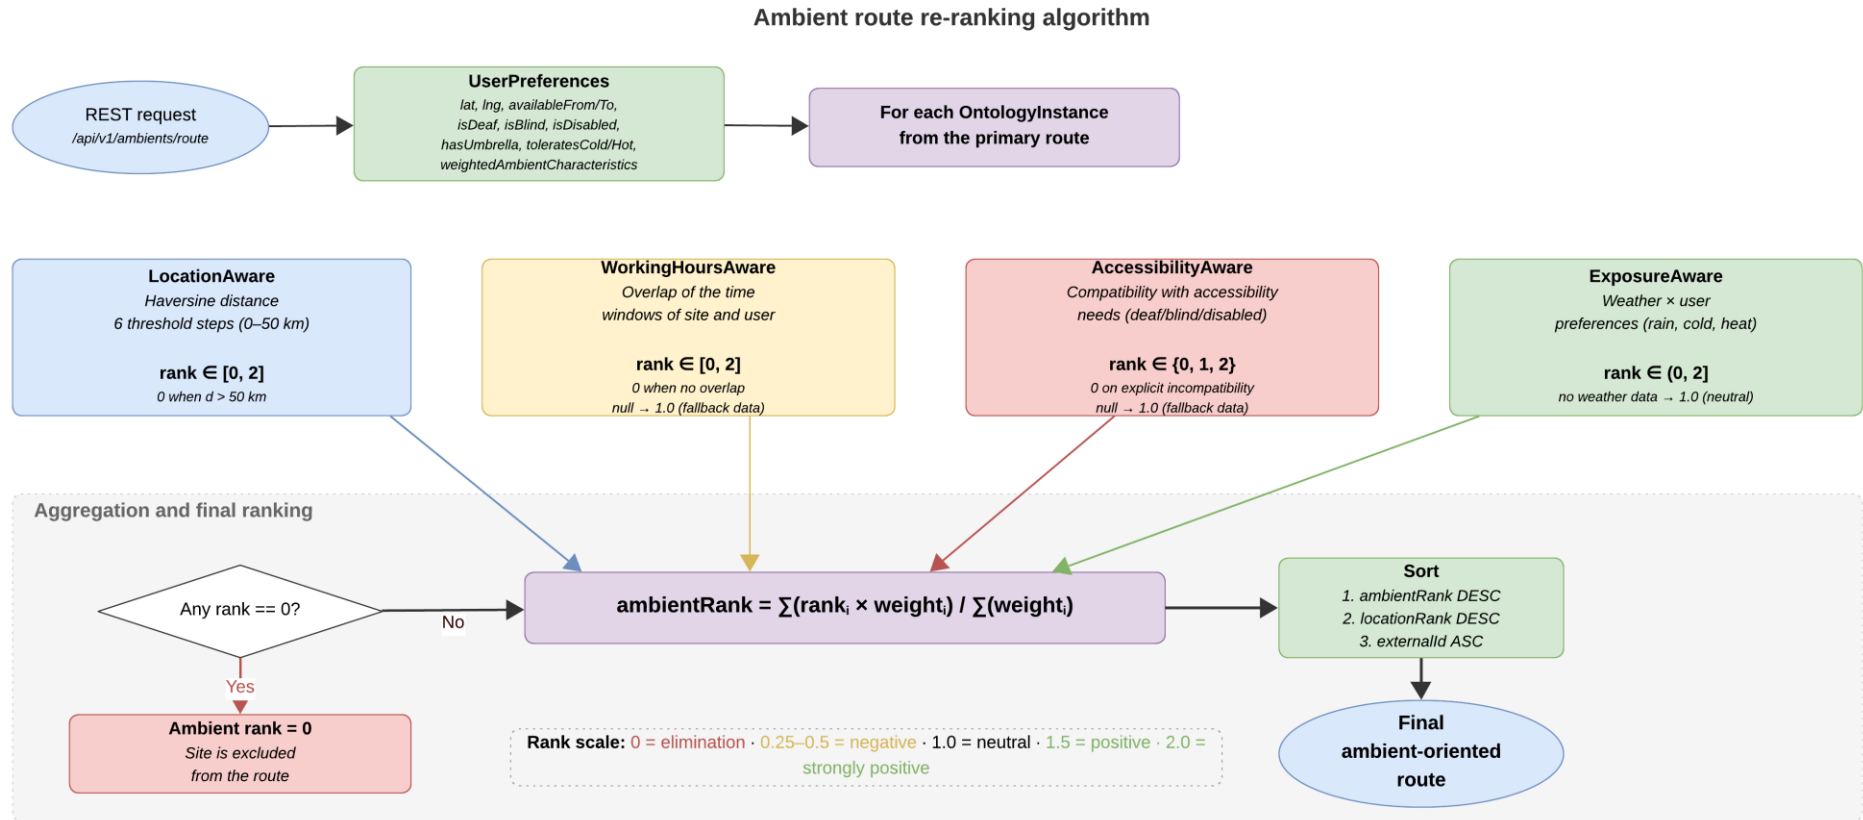


Figure 3. Flowchart of the algorithm for ambient re-evaluation and ranking of candidate sites using personalized information and an initial list of ontological instances.

The algorithm provides explainability by construction: for each site, the intermediate scores per factor and their weights are available, the elimination logic is explicit, and the DTO models include the contextual data that participated in the evaluation — for example, whether it is raining, whether it is hot or cold.

The implemented endpoint `GET /api/v1/ambients/{id}/explain` returns `ExplanationDTO` with an explicit classification of the effect of each factor (`ELIMINATED`, `STRONGLY_NEGATIVE`, `NEGATIVE`, `NEUTRAL`, `POSITIVE`, `STRONGLY_POSITIVE`), a textual reason and a summarized description. For each factor the rank, the weight, whether the value is inferred or fallback, and a textual explanation of its specific contribution are provided. The explanation is constructed deterministically from the data that the algorithm already computes, without an additional inference layer.

Uncertainty is treated conservatively: missing values receive a neutral rank (1.0), while an explicitly established incompatibility in accessibility or working hours leads to elimination. The system distinguishes directly available, inferred and fallback values, which allows the explanation layer to indicate the origin of each factor.

3.3 Limitations of the prototype

The prototype has four clearly recognizable limitations: (1) no full route optimization with transport times and a global daily programme — the focus is on re-ranking candidates; (2) the database contains synthetically generated users and ratings, which limits behavioural inferences; (3) AI inference uses a probabilistic model that does not replace expert verification as a source of truth; (4) the REST interface is sufficient for the prototype but not finished as a public industrial API, as administrator operations and a richer classification of errors are missing.

3.4 Toolkit for expert validation

For expert validation of the prototype, a dedicated checklist has been developed for reviewing newly added or AI-enriched sites, covering semantic description and ontological categorization, spatial data, working hours, multidimensional accessibility, exposure and recommendation explanation. The full toolkit and its application procedure are described in the dissertation and turn expert evaluation from an intuitive activity into a more structured and repeatable procedure, useful both for operational support and for future experiments and for demonstrating the reliability of the system before a scientific audience.

3.5 Analysis of the influence of individual ambient factors

The ambient factors carry different types of constraint: location is a spatial filter, working hours is a temporal filter, accessibility is a personal filter tied to route feasibility, and exposure is a contextual adjustment factor for comfort and reasonableness of the visit under specific weather conditions.

The exposure rank is formed multiplicatively from an initial value of 1.0: under precipitation without shelter the rank is divided by 4, under cold without tolerance — by 4, and under favourable weather or the presence of shelter/umbrella — multiplied by 2 (capped at 2.0 via $\min(1.5 \times \text{rank}, 2.0)$ to avoid overcompensation). The multiplicative nature, combined with the absence of a hard-elimination condition, explains why exposure is a corrector, not a filter: it does not remove sites, but can substantially change their ordering.

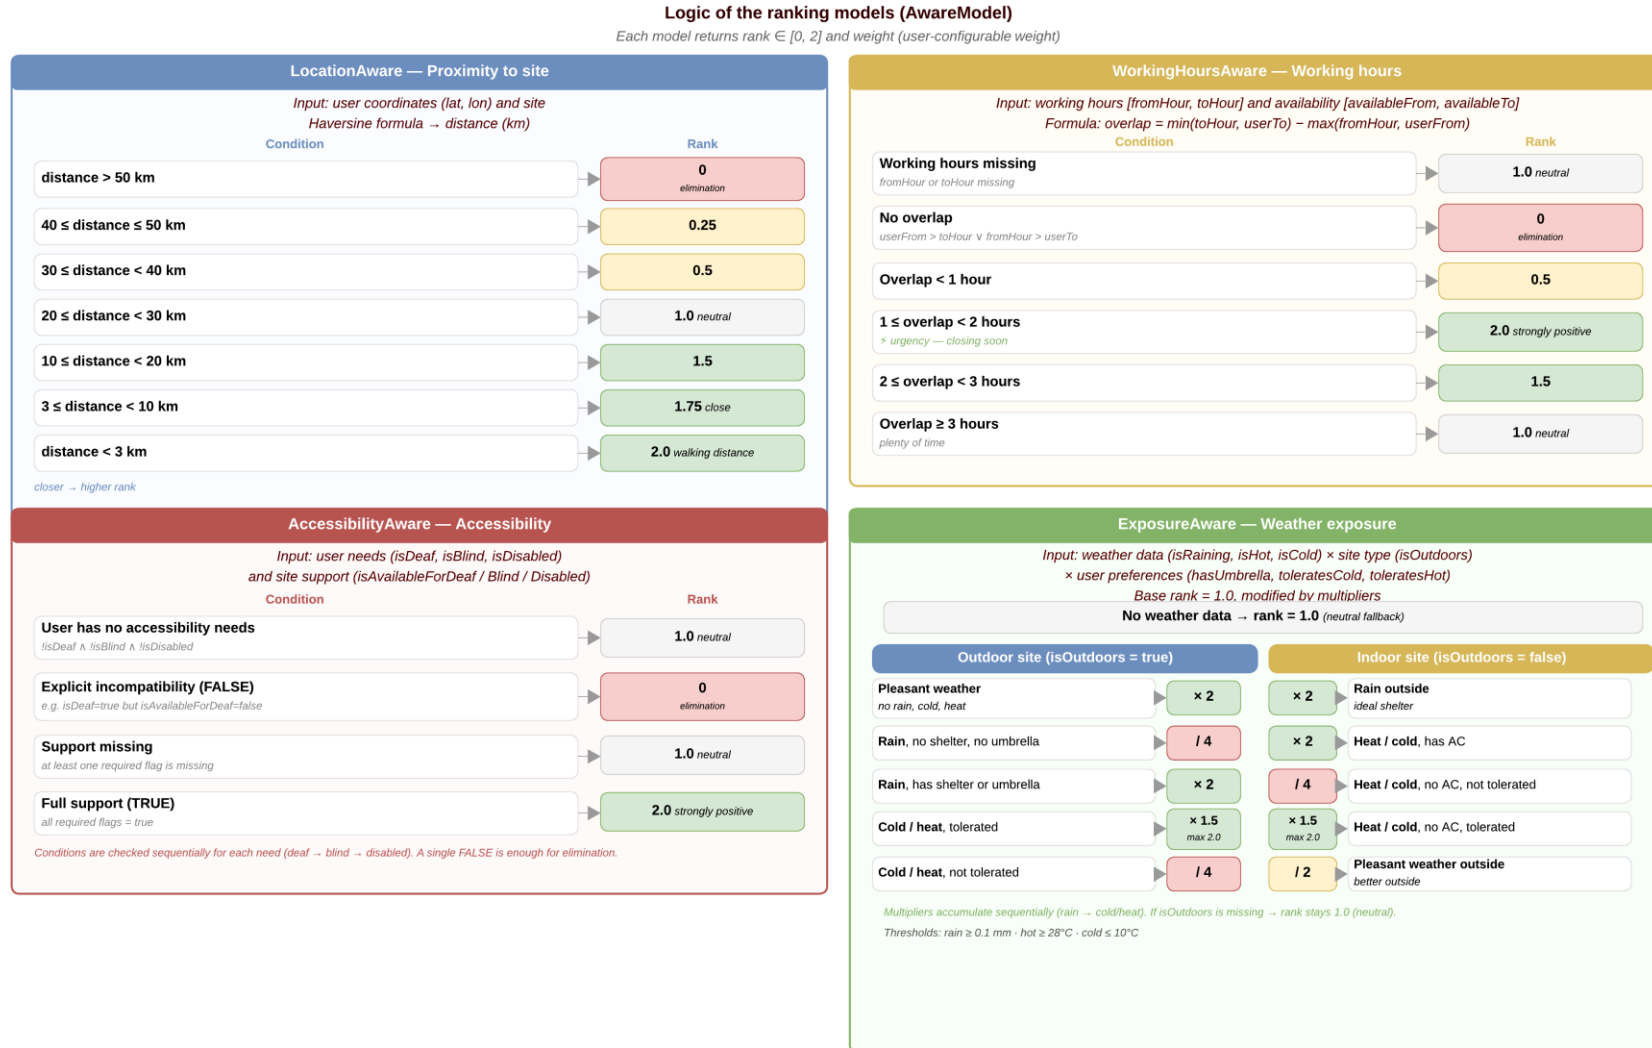


Figure 4. Decision tables describing the ranking logic of the individual AmbientCharacteristicAwareModel implementations.

The absence of feedback data in the "cold start" situation justifies the choice of a heuristic rather than a fully trainable model at the present stage: heuristically formalized factors are directly interpretable and defensible, without precluding future machine learning — an intermediate, yet strongly explainable, solution.

The process of ambient re-ranking is based on a mathematical model for a weighted assessment of the applicability of each tourist site, the formal description of which follows.

Let $C = \{c_1, c_2, \dots, c_m\}$ be the set of ambient characteristics (e.g. weather conditions, accessibility, time window), and $W = \{w_1, w_2, \dots, w_m\}$ their corresponding weights, defined through the user profile. The final applicability score $A(O)$ for a given site O is computed by the formula:

$$A(O) = \frac{\sum_{j=1}^m r(c_j, O) \cdot w_j}{\sum_{j=1}^m w_j}$$

where $r(c_j, O) \in [0, 2]$ represents the coefficient of correspondence of the site with respect to the specific characteristic c_j ; a value of 0 denotes critical incompatibility and leads to the elimination of the site from the route.

Algorithm 1. Ambient ranking of a site

Input: ontological instance O , set of characteristics $C(O)$ and their corresponding weights W .

Output: final score $A(O) \in (0, 2]$ or 0 (elimination)

```

1. sumRatings  $\leftarrow$  0, sumWeights  $\leftarrow$  0
2. For each characteristic  $c_i \in C(O)$ :
3.    $r_i \leftarrow r(c_i, O)$ 
4.    $w_i \leftarrow W[c_i]$ 
5.   If  $r_i = 0$ : return 0 // critical incompatibility
6.   sumRatings  $\leftarrow$  sumRatings +  $r_i \times w_i$ 
7.   sumWeights  $\leftarrow$  sumWeights +  $w_i$ 
8.  $A(O) \leftarrow$  sumRatings / sumWeights
9. Return  $A(O)$ 

```

If the input initial route is empty, the system retrieves via spatial query sites within a 50 km radius and applies the above algorithm per candidate; results are sorted descending by final rank with complexity $O(n \cdot m)$. The weighted arithmetic mean with hard elimination at zero rank separates soft preferences from hard constraints and provides an interpretable baseline model, where n is the number of candidates and m the number of ambient characteristics.

Table 3. Ambient factors in the re-ranking algorithm.

Factor	Input data	Thresholds / scale	Effect
Location	coordinates of user and site	0–3 km $\rightarrow$ 2.0; 3–10 $\rightarrow$ 1.75; 10–20 $\rightarrow$ 1.5; 20–30 $\rightarrow$ 1.0; 30–40 $\rightarrow$ 0.5; 40–50 $\rightarrow$ 0.25; >50 $\rightarrow$ 0	>50 km $\rightarrow$ elimination; threshold heuristic
Working hours	access hours, free window	no overlap $\rightarrow$ 0; <1 h $\rightarrow$ 0.5; 1–2 h $\rightarrow$ 2.0; 2–3 h $\rightarrow$ 1.5; >3 h $\rightarrow$ 1.0	no overlap $\rightarrow$ elimination; max rank at 1–2 h
Accessibility	motor, visual, auditory	no constraint $\rightarrow$ 1.0; compatible $\rightarrow$ 2.0; incompatible $\rightarrow$ 0; unknown $\rightarrow$ 1.0	incompatibility $\rightarrow$ elimination; unknown $\rightarrow$ neutral
Exposure	indoor/outdoor, shelter, air conditioning, weather	multiplicatively from 1.0; rain without shelter $\div$ 4; cold without tolerance $\div$ 4; favourable weather $\times$ 2; max 2.0	corrector, not filter; demotes or boosts, does not eliminate

3.6 Extensions of the data and knowledge model

The current data model is sufficient for the implemented prototype, but leaves natural points for extension: adding new ambient characteristics (price, seasonality, congestion, transport accessibility, terrain difficulty) without radical architectural changes; enriching the ontology layer with thematic routes by historical period or institutional affiliation; and formally recording data provenance and confidence — especially useful for a stricter scientific interpretation of uncertainty.

3.7 Ethical and applied aspects of the proposed solution

Working with accessibility, location and personal constraints requires ethical decisions at the prototype level: passwords are stored hashed, sessions are stateless through JWT with a limited validity period, and the user's coordinates are stored as part of the profile and are not refreshed automatically — the system does not support background tracking and does not store trajectories. Architectural decisions on security are thus not only technical, but also ethical.

The ethical aspect is also linked to AI inference. If the language model is used without control, there is a risk that the system may “invent” properties of the site and mislead the user. For this reason the present work adopts a conservative position: AI inference is used only as a mechanism for structured enrichment of the data and does not replace domain verification. This is a more responsible setup than uncritically trusting a probabilistic model and is consistent with the principles of Regulation (EU) 2024/1689 on artificial intelligence [54].

3.8 Conclusions of Chapter 3

Several important conclusions follow from the analysis of the implementation.

- The proposed model is technologically achievable through the use of contemporary server technologies, a standard REST interface and a relational database.
- Ontological search and ambient evaluation can be integrated sequentially without breaking the modularity of the architecture.
- The user profile is sufficiently rich to support working personalization from the very first use.
- The weather context and AI inference can be included as external sources of knowledge that strengthen the system, provided that they are managed through cautious fallback mechanisms.
- The REST interface of the system materializes clearly delineated domain responsibilities and provides a good basis for future extension.

In the next chapter, these architectural and algorithmic decisions will be subjected to experimental validation on a synthetic and a real corpus of data.

Chapter 4. Experimental validation of the proposed model

4.1 Experimental setup

The experimental validation uses two complementary corpora: a controlled synthetic set of 100 sites in Plovdiv and a real corpus of 143 sites from across Bulgaria. The two corpora cover mutually exclusive risks: the synthetic one allows clean ablation experiments under full expert control of the distribution, while the real one tests the robustness of AI inference against authentic missing values. Results are presented in parallel, so that every quantitative claim is verified under both a controlled and a realistic setup across hard constraints and various user profiles.

Within this setup a controlled corpus with expert-defined ambient characteristics was constructed and three strategies were compared: proximity without personalization, ambient evaluation without dynamic weather context and the proposed ambient model. The methodology was supplemented with ablation experiments and scenario-based profile validation, so that the quantitative analysis takes into account the real effect of location, accessibility, exposure and working hours in the re-ranking process.

Evaluation is carried out for five profiles: a standard tourist, a wheelchair user, a tourist on a rainy day without an umbrella, a tourist on a hot day without tolerance for high temperature and a user with a limited time window. For each profile an expert set of relevant sites is defined, which allows comparison of the top 10 results on a catalogue in which the recommended 10 positions account for only 10% of all available candidates.

The quality of the ranking is evaluated with Precision@k and nDCG@k [55], supplemented with the Feasibility % metric — the share of the returned top-10 recommendations that retain a positive rank after the full ambient model is applied. The “working hours” factor modulates the rank along a non-monotonic urgency curve: with overlap < 1 h the site receives a rank of 0.5 (a risky but admissible opportunity); < 2 h — 2.0 (maximum urgency); < 3 h — 1.5 (moderate acceleration); ≥ 3 h — 1.0 (neutral). The curve directs attention to time-pressed sites, rather than linearly to a longer overlap.

4.2 Quantitative results from the comparison of strategies

Table 4 summarizes the comparison of the three strategies on the synthetic corpus, and Table 5 (§4.6) — on the real corpus. The k values differ: the synthetic corpus is evaluated at k = 5 and k = 10, while the real corpus at k = 3 and k = 5, since after spatial filtering the surviving candidates in the radius around Gabrovo are 14–25, which would make k = 10 trivially padded.

For the standard profile the proximity baseline reaches $P@5 = 0.400$ and $nDCG@10 = 0.376$, while the full ambient model reaches $P@5 = 1.000$ and $nDCG@10 = 0.857$ thanks to the preference for outdoor sites under pleasant weather. The strongest effect is for the wheelchair profile (baseline $P@10 = 0.100 \rightarrow$ full model 0.900; feasibility from 70% to 100% after elimination of 25 inaccessible sites and 7 with incompatible working hours). Under rain and heat the "Ambient without weather data" strategy matches the baseline — confirming that it is the dynamic weather layer that brings the added value.

Under a limited time window (16:00–18:00), the full ambient model reaches $P@5 = 1.000$, $P@10 = 0.700$ and $nDCG@10 = 0.791$ against 0.237 for the baseline ($\Delta = +0.554$). The difference is due to the combined effect of the exposure factor (preference for outdoor sites) and the urgency model for working hours: sites with 1 h overlap receive the maximum urgency rank (2.0), while sites with full coverage (2 h) receive a moderate rank (1.5), prioritizing sites that will close soon.

4.3 Ablation analysis of the ambient factors

In order to assess the effect of each ambient factor in isolation, an ablation analysis was carried out, in which the individual characteristics are selectively turned off in the model without changing the rest of the architecture. The analysis is carried out in parallel on the two corpora: Table 6 presents the summarized results for the two corpora. The comparison between the two ablation matrices is discussed in Sec. 4.6.

The ablation analysis uses a full $nDCG@10$ matrix under selective removal of each ambient factor.

For the standard profile, removing exposure brings $nDCG@10$ down from 0.857 to 0.376 ($\Delta = +0.482$) — distinguishing outdoor from indoor sites is the main driver; location shows no ablation difference due to the dense intra-urban distribution, in which all sites fall into two adjacent distance classes without sufficient dispersion.

For the wheelchair user profile, exposure yields the greatest value ($\Delta = +0.655$), followed by accessibility ($\Delta = +0.422$) — accessibility acts as a filter eliminating 25 inaccessible sites and 7 with incompatible working hours, while exposure provides granular ranking among the 68 surviving candidates. Under rain and heat, the "Ambient without weather data" strategy matches the proximity baseline (0.763 and 0.624), confirming that the dynamic weather context is a necessary component of the ambient model.

4.4 Scenario validation and qualitative analysis

A scenario validation checks the domain invariants (elimination of infeasible sites, promotion of suitable candidates, stability under incomplete data) on both corpora; Table 7 (§4.6) summarizes surviving and eliminated sites and key metrics for each profile across the two corpora.

The larger number of eliminated sites (32 against the originally introduced 25 inaccessible) is due to the fact that the ambient filter is cumulative — in addition to the physical barriers, the system also eliminates in parallel sites that are temporally incompatible with the user's profile.

4.5 Discussion of results and limitations of the validation

The combined analysis shows that the model is a context-aware re-ranking layer, not a universal replacement for the proximity baseline. Standard profile: $P@5 = 1.000$ vs 0.400; with additional constraints — up to a ninefold increase in $P@10$ for the wheelchair profile and +50% in $nDCG@10$ on a hot day. The synthetic corpus is expertly constructed and the baseline is deliberately elementary — the profiles are purely ambient, with no thematic component against which a semantic or collaborative baseline would have anything to compare.

4.6 Validation with real data from the subject domain

In order to complement the validation on the synthetic corpus, the present section presents the results of applying the same set of strategies and metrics to real data from the subject domain. The corpus includes 143 sites of Bulgaria's cultural and historical heritage, loaded from the production JSON file of the prototype (`raw-ontology-instances.json`) and transformed into `OntologyInstance` objects by a test execution task, which reproduces the structure expected by the ranking layer. The primary source for the names, descriptions, coordinates and initial categorization, from which the ontological concepts are derived, is the public Trips Journal catalogue, "Sights" section, from which the data were collected via a script for the purposes of the prototype [56]. Unlike the production import, the test process reads pre-cached AI-inferred characteristics in order to guarantee reproducibility without API calls on every benchmark run. The centre of the experiment is Gabrovo (42.8742°N, 25.3187°E), chosen because of the presence of a sufficient number of accessible sites in the vicinity for validating the user-with-disabilities profile.

Origin of the data: the coordinates of the 143 sites are real (from the primary JSON corpus); working hours, accessibility and exposure are inferred via an AI model (GPT-5.4 at temperature = 0.0) with prompts functionally equivalent to the production services. API calls for accessibility and exposure succeeded for all 143 sites (0% errors); for working hours 142/143 (99.3%). The share of unresolved fields ranges from 7.7% (`isOutdoors`) to 74.8% (`isAvailableForDeaf`); for wheelchair accessibility it is 54.5%. Unresolved fields receive a neutral rank (1.0) uniformly across all three characteristics.

An expert correction of accessibility was simulated: for 54.5% of sites the `isAvailableForDisabled` field remains unresolved, and for indoor museums a conservative expert correction was applied in which wheelchair accessibility is assumed present when no indicators of inaccessibility are available. The step is marked as simulated and serves as a sensitivity check, rather than as a replacement for actual expert annotation.

On the real data, the full model reaches $nDCG@5 = 1.000$ for three of the five profiles (standard, wheelchair, limited window) and $P@5 = 1.000$ for the wheelchair profile vs 0.600 (feasibility 60%) for the baseline; under a hot day the improvement is +31% in $nDCG@5$. Ablation confirms the factor profile from the synthetic experiment: accessibility and location under mobility constraints, exposure on a hot day; under rain the deltas are zero because of the specific geographical configuration around Gabrovo, in which the nearest sites have shelter.

The two corpora — synthetic and real — complement each other: the synthetic enables clean ablation conclusions, while the real confirms that the effects reproduce under a less controlled distribution. The whole setup is fully reproducible without external API calls: the synthetic corpus is generated programmatically, the real AI inferences are cached in the source code of the prototype, and the two benchmarks run through the project’s standard test tasks (`BenchmarkOutputTest` and `RealDataBenchmarkOutputTest`).

4.7 Conclusions of Chapter 4

The systematic verification shows that the ambient-oriented model brings added value across all tested scenarios — the strongest effect for profiles sensitive to accessibility and weather conditions (up to ninefold improvement in $P@10$ for the wheelchair profile, from 0.100 to 0.900), with 100% feasibility guaranteed by construction.

Since the reference sets are built by the author from explicit domain rules rather than through independent expert annotation, the results should be regarded as a systematic check of internal consistency and a proof of concept. Outlined directions for development: more realistic corpora, externally validated data, stronger baselines, verification with real users, sensitivity analysis of the chosen thresholds and an empirical check of the non-monotonic urgency curve against simpler monotonic alternatives.

Table 4. Comparison of strategies on a synthetic corpus of 100 sites.

Profile	Proximity only $nDCG@10$	Without weather $nDCG@10$	Full model $nDCG@10$	Full model $P@5$
Standard	0.376	0.376	0.857	1.000
Wheelchair	0.139	0.272	0.927	1.000
Rainy day	0.763	0.763	0.870	1.000
Hot day	0.624	0.624	0.934	1.000
Limited time	0.237	0.511	0.791	1.000

Table 5. Comparison of strategies on a real corpus of 143 sites.

Profile	Proximity only nDCG@5	Without weather nDCG@5	Full model nDCG@5	Full model P@5
Standard	0.869	0.869	1.000	1.000
Wheelchair	0.723	1.000	1.000	1.000
Rainy day	0.684	0.684	0.684	0.600
Hot day	0.553	0.553	0.723	0.600
Limited time	0.869	0.830	1.000	1.000

Table 6. Summary of the ablation analysis.

Profile	Synthetic: largest drop	Real: largest drop
Standard	Exposure +0.482	Exposure +0.131
Wheelchair	Exposure +0.655; accessibility +0.422	Accessibility +0.661; location +0.384
Rainy day	Exposure +0.107	no distinct drop
Hot day	Exposure +0.309	Exposure +0.170
Limited time	Exposure +0.280; working hours +0.077	Location +1.000

Table 7. Summary of the scenario validation: surviving and eliminated sites by profile and corpus.

Profile	Synth.: Total	Synth.: Surviving / Elim.	Synth.: nDCG@10	Real: Total	Real: Surviving / Elim.	Real: nDCG@5
Standard tourist	100	99 / 1	0.857	143	25 / 118	1.000
Wheelchair	100	68 / 32	0.927	143	14 / 129	1.000
Rainy day without umbrella	100	99 / 1	0.870	143	25 / 118	0.684
Hot day	100	99 / 1	0.934	143	25 / 118	0.723
Limited time window	100	99 / 1	0.791	143	25 / 118	1.000

Conclusion

The present work proposes and substantiates a model of an ambient-oriented intelligent tourist guide that combines an ontological representation of cultural and historical sites, an explicit user profile with weights, a dynamic external context and a cautious mechanism for completing missing ambient data. The main thesis: a tourist recommendation is not only a problem of retrieving relevant sites, but also of assessing their applicability in a specific situation.

The literature analysis outlines a gap between semantic tourist guides, context-aware systems, accessibility models and the new AI approaches for data enrichment; the result of the work is in a unified architectural and algorithmic setup that links these lines into a common model, instead of regarding them as isolated techniques.

At the conceptual level, the work formalizes the role of ambient factors as constraints and correctors with different semantics over the ranking. Particularly significant are the multidimensional representation of accessibility, the separation of exposure as a contextual factor, and the treatment of uncertainty as a manageable property of the system.

At the applied level, a working prototype has been implemented (REST, ontology and ambient layers, relational database, Open-Meteo, controlled AI inference) with pre-defined scenario conditions for reproducibility. The ablation analysis over the two corpora shows an individual measurable effect: location dominates in spatial scenarios, accessibility — under motor constraints, exposure — under meteorological stress, and working hours — under a restricted time window.

On the synthetic corpus of 100 sites, the full ambient model raises $nDCG@10$ for the standard profile from 0.376 to 0.857, under heat from 0.624 to 0.934, and under a restricted window from 0.237 to 0.791; for the wheelchair profile $P@10$ rises from 0.100 to 0.900, and feasibility — from 70% to 100%.

On the real corpus of 143 sites the full model reaches $nDCG@5 = 1.000$ for three of the five profiles — a result of the strength of the elimination filter (fixed location Gabrovo shrinks the set to 14–25 surviving sites). A more substantial indicator is the behaviour on the wheelchair profile and under meteorological stress, where the ambient factors make the difference with respect to the proximity strategy.

Limitations: the reference sets are built from objective domain rules (a check of internal consistency); the methodology is applicable with an external reference set without changing the setup. The high share of eliminated sites on the real corpus (over 90%) limits the statistical power over the small surviving set. The prototype does not implement global route optimization — the scoring layer works on list ordering, rather than over a traversal with transport times.

In summary, the dissertation proposes and tests a framework in which ontological retrieval, multidimensional accessibility, weather-conditioned exposure, a parametrizable profile, factor-based explanation and controlled AI completion are connected in a common model. The result lies in the semantics of the factors and their interaction; the prototype demonstrates each component's own measurable contribution and is designed for extension toward richer ambient factors and full route optimization and user validation.

Results of the dissertation

Conceptual and scientific-applied results

- A two-stage architectural setup: semantic retrieval of candidates followed by ambient re-ranking through a set of ambient characteristics with distinct semantic functions (spatial, temporal, personal, contextual) — a basis for ranking and for explanation of the obtained recommendations.

- A multidimensional model of accessibility (motor, visual, auditory) treated as real constraints on the ranking, integrated with a dynamic weather context as an external factor in exposure assessment — the model responds to personal and environmental conditions, not only to static properties of the site.
- A cautious mechanism has been implemented for completing missing ambient data through a language model (currently GPT-5.4) with structured outputs (JSON schema), domain rules and distinction of explicit values from inferred ones. Unresolved fields receive a neutral rank (1.0), consistently across accessibility, exposure and working hours, and the architecture is open to subsequent expert verification through a data-provenance and confidence layer. The result lies in the controlled use of AI, not in the automatic acceptance of the inferred values as fact.

Applied results

- A modular server-side prototype has been implemented with a REST interface, ontology layer, relational database with JOINED inheritance, a parametrizable user profile, Open-Meteo integration, controlled AI inference with circuit breaker and bulkhead protection, factor-based explanation and storage of user ratings (intended as a basis for future hybrid recommendation).

Methodological results

- A methodology for validating ambient-oriented route recommendations through comparative baseline strategies, ablation analysis and scenario-based profile checking on a controlled corpus; it can be applied with an external reference set without changing the setup, and the benchmarks are implemented as automated tests embedded in the prototype, which guarantees reproducibility.
- Explainability, accessibility and uncertainty management are argued to be first-class properties of the intelligent tourist guide, rather than as peripheral additions.

Directions for future research

A natural direction for development is the transition from ambient re-ranking to multi-criteria route planning (transport times, visit sequence, duration of stay, overall utility), in which the ambient model is used as an evaluation layer within a more complex optimization mechanism.

Beyond route optimization, the ambient layer can be enriched with new characteristics — price, seasonality, congestion, safety, transport accessibility, terrain difficulty, proximity to parking, food or sanitary services — which makes the route even more sensitive to the real conditions of the visit.

A more rigorous handling of uncertainty includes explicit storage of data provenance, AI-inference confidence and the date of the last verification, which would allow the data to be arranged in a hierarchy according to the degree of trust.

The accumulated user ratings open the possibility of hybrid recommendation, in which behavioural signals enrich the profile-contextual model without replacing its explainability; the factor-based explanation layer can be augmented with natural-language formulations, visualizations and interactive exploration of alternatives, so the user receives both the final result and an understandable rationale for the ordering.

As a natural continuation of the directions outlined, several research questions and working hypotheses can be formulated, which follow from the present dissertation. They show that the proposed model is not a closed construction, but a starting point for a next stage of systematic scientific research.

- Working hypothesis 1: the inclusion of a multidimensional accessibility model as a first-class factor raises the practical applicability of routes compared with models that use only proximity and topic. This hypothesis can be verified by comparing the share of feasible routes for profiles with constraints with and without an accessibility layer, where the key measure is the share of routes with at least 5 applicable sites.
- Working hypothesis 2: the inclusion of an exposure model sensitive to meteorological context improves recommendation quality under adverse conditions. Verification of this hypothesis can be performed through scenario experiments under precipitation, high temperature and low temperature, comparing the ambient-oriented strategy with a version without meteorological context, with necessary supplementation by a user study rather than IR metrics alone.
- Working hypothesis 3: controlled AI inference of missing ambient characteristics leads to higher coverage of applicable routes compared with a regime in which missing values are ignored. The key balance here is between coverage and reliability. For this reason, the hypothesis should be verified not only through quantitative metrics, but also through expert assessment of the plausibility of the inferred values, with the central measure being the accuracy of the inferred values against expert annotation.
- Working hypothesis 4: the presence of factor explanations — increasingly relevant in the context of the regulatory requirements for transparency (Regulation (EU) 2024/1689, EU AI Act [54]) — raises users' trust in the system and their willingness to follow the recommendations. Such a hypothesis can be verified through a user study with control groups, in which one group receives only a list of recommendations and another — the same recommendations with brief explanations of the reasons for the ordering.
- Working hypothesis 5: explicit profiling is more effective at the early stage of use than models relying only on accumulated behavioural signals. This is particularly relevant in the tourist domain, where user behaviour is often episodic. Verification of this hypothesis requires a comparison between an explicit profile, a behavioural model and a combined approach for new users.

Two research questions remain open: to what extent the weights of ambient factors can be automatically calibrated rather than fully user-provided (semi-adaptive models with explicit initialization and behavioural calibration), and to what extent the same ambient model can be transferred to other domains — urban services, health tourism, educational routes or ecological walks. The outlined hypotheses and questions sketch a concrete research program on the infrastructure of the current prototype.

Publications on the topic

The main publications on which the present dissertation rests and which it develops are the following:

1. **Georgi Kostadinov**; Kristian Milev; Stefan Staynov; Asya Stoyanova-Doycheva, Algorithm for Generating and Visualizing Routes of an Intelligent Tourist Guide, Automatics and Informatics (ICAI), 1-3 Oct. 2020, Varna, Bulgaria, Article number 9311310, DOI: 10.1109/ICAI50593.2020.9311310, (IEEE, SCOPUS)
 - <https://ieeexplore.ieee.org/document/9311310>
2. Stefan Staynov, Kristian Milev, **Georgi Kostadinov**, Asya Stoyanova-Doycheva, Traditional Bulgarian Costume Classification in a Tourist Guide application, Automatics and Informatics (ICAI), 1-3 Oct. 2020, Varna, Bulgaria, Article number 9311384, DOI: 10.1109/ICAI50593.2020.9311384, (IEEE, SCOPUS)
 - <https://ieeexplore.ieee.org/document/9311384>
3. Kristian Milev, **Georgi Kostadinov**, Stefan Staynov, Asya Stoyanova-Doycheva, Intelligent Tourist Guide – Generating Ambient-Oriented Routes and Digital Classification of Elements

of the Cultural and Historical Heritage of Bulgaria, Cultural and Historical Heritage: Preservation, Presentation, Digitalization, KIN Journal, volume 7, Issue 2, 2021, ISSN: 2367-8038, DOI: doi.org/10.26615/issn.2367-8038.2021_2_006

○ http://www.math.bas.bg/vt/kin/files/papers/7_2/06-KIN-7-2-2021.pdf

4. **Georgi Kostadinov**, Kristian Milev, Stefan Staynov, Asya Stoyanova-Doycheva, Applying Ambient Characteristics, Personalization and Ontology Concepts to Tourist Guide Route Generation, Vol. 12 (2022): Digital Presentation and Preservation of Cultural and Scientific Heritage, pp. 229-234, 2022, ISSN: 1314-4006 (Print), eISSN: 2535-0366 (Online), (WoS, SCOPUS)

○ <https://doi.org/10.55630/dipp.2022.12.19>

○ <https://dipp.math.bas.bg/dipp/article/view/dipp.2022.12.19/pdf>

○ <https://www.webofscience.com/wos/author/record/1390446,30412672>

References

- [1] D. Gavalas, C. Konstantopoulos, K. Mastakas and G. Pantziou, "Mobile Recommender Systems in Tourism," *Journal of Network and Computer Applications*, vol. 39, pp. 319-333, 2014.
- [2] S. Halder, K. H. Lim, J. Chan and X. Zhang, "A survey on personalized itinerary recommendation: From optimisation to deep learning," *Applied Soft Computing*, vol. 152, p. 111200, 2024.
- [3] A. Stoyanova-Doycheva, T. Glushkova, V. Ivanova, L. Doukovska and S. Stoyanov, "A Multi-agent Environment Acting as a Personal Tourist Guide," in *Castillo, O., Melin, P., Kacprzyk, J. (eds) Intuitionistic and Type-2 Fuzzy Logic Enhancements in Neural and Optimization Algorithms: Theory and Applications. Studies in Computational Intelligence*, vol. 862, Springer, Cham, 2020.
- [4] M. Miteva, A. Stoyanova-Doycheva and V. Ivanova, "Architecture of an Intelligent Tourist Guide," *Computer Sciences and Communications*, vol. 7, no. 1, 2018.
- [5] T. Glushkova, M. Miteva, A. Stoyanova-Doycheva, V. Ivanova and S. Stoyanov, "Implementation of a Personal Internet of Thing Tourist Guide," *American Journal of Computation, Communication and Control*, 2018.
- [6] G. Kostadinov, K. Milev, S. Staynov and A. Stoyanova-Doycheva, "Applying Ambient Characteristics, Personalization and Ontology Concepts to Tourist Guide Route Generation," in *Digital Presentation and Preservation of Cultural and Scientific Heritage (DiPP)*, 2022.
- [7] G. Kostadinov, K. Milev, S. Staynov and A. Stoyanova-Doycheva, "Algorithm for Generating and Visualizing Routes of an Intelligent Tourist Guide," in *International Conference Automatics and Informatics (ICAI)*, Varna, Bulgaria, 2020.
- [8] G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper and M. Pinkerton, "Cyberguide: A Mobile Context-Aware Tour Guide," *Wireless Networks*, vol. 3, no. 5, pp. 421-433, 1997.
- [9] A. Garcia, O. Arbelaitz, M. T. Linaza, P. Vansteenwegen and W. Souffriau, "Personalized Tourist Route Generation," in *Daniel, F., Facca, F. M. (eds) Current Trends in Web Engineering. ICWE 2010. Lecture Notes in Computer Science*, vol. 6385, Springer, Berlin, Heidelberg, 2010.
- [10] N. Davies, K. Cheverst, K. Mitchell and C. Efstratiou, "Using and Determining Location in a Context-sensitive Tour Guide," *IEEE Computer*, 2001.
- [11] G. Adomavicius and A. Tuzhilin, "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734-749, 2005.
- [12] L. Wu, Z. Zheng, Z. Qiu, H. Wang, H. Gu, T. Shen, C. Qin, C. Zhu, H. Zhu, Q. Liu, H. Xiong and E. Chen, "A Survey on Large Language Models for Recommendation," *World Wide Web*, 2024.
- [13] G. Adomavicius, B. Mobasher, F. Ricci and A. Tuzhilin, "Context-Aware Recommender Systems," *AI Magazine*, vol. 32, no. 3, pp. 67-80, 2011.
- [14] R. Burke, "Hybrid Recommender Systems: Survey and Experiments," *User Modeling and User-Adapted Interaction*, vol. 12, no. 4, pp. 331-370, 2002.
- [15] P. Lops, M. de Gemmis and G. Semeraro, "Content-based Recommender Systems: State of the Art and Trends," in *Ricci, F., Rokach, L., Shapira, B., Kantor, P. B. (eds) Recommender Systems Handbook*, Springer, 2011.
- [16] R. Kramer, M. Modsching and K. ten Hagen, "Development and Evaluation of a Context-Driven, Mobile Tourist Guide," *International Journal of Pervasive Computing and Communications*, vol. 3, no. 4, pp. 378-399, 2007.
- [17] C. Lamsfus, A. Alzua-Sorzabal, D. Martin, Z. Salvador and A. Usandizaga, "Human-centric Ontology-based Context Modelling in Tourism," in *Proceedings of KEOD*, 2009.
- [18] A. Sieg, B. Mobasher and R. Burke, "Learning Ontology-based User Profiles: A Semantic Approach to Personalized Web Search," *IEEE Intelligent Informatics Bulletin*, 2007.
- [19] A. García-Crespo, J. L. López-Cuadrado, R. Colomo-Palacios, I. González-Carrasco and B. Ruiz-Mezcua, "Sem-Fit: A Semantic Based Expert System to Provide Recommendations in the Tourism Domain," *Expert Systems with Applications*, 2011.
- [20] T. Voelkel and G. Weber, "Routecheckr: Personalized Multicriteria Routing for Mobility Impaired Pedestrians," in *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '08)*, 2008.
- [21] N. Tintarev and J. Masthoff, "Designing and Evaluating Explanations for Recommender Systems," in *Recommender Systems Handbook*, Springer, 2011.
- [22] Y. Zhang and X. Chen, "Explainable Recommendation: A Survey and New Perspectives," *Foundations and Trends in Information Retrieval*, vol. 14, no. 1, pp. 1-101, 2020.
- [23] S. Darcy and T. Dickson, "A Whole-of-Life Approach to Tourism: The Case for Accessible Tourism Experiences," *Journal of Hospitality and Tourism Management*, 2009.

- [24] D. Buhalis and S. Darcy, *Accessible Tourism: Concepts and Issues*, Channel View Publications, 2011.
- [25] F. Ricci, "Travel Recommender Systems," *IEEE Intelligent Systems*, 2002.
- [26] W. Souffriau, P. Vansteenwegen, J. Vertommen, G. Vanden Berghe and D. Van Oudheusden, "A Personalized Tourist Trip Design Algorithm for Mobile Tourist Guides," *Applied Artificial Intelligence*, vol. 22, no. 10, 2008.
- [27] P. Vansteenwegen, W. Souffriau and D. Van Oudheusden, "The Orienteering Problem: A Survey," *European Journal of Operational Research*, vol. 209, no. 1, pp. 1-10, 2011.
- [28] L. Ardissono, A. Goy, G. Petrone, M. Segnan and P. Torasso, "Intrigue: Personalized Recommendation of Tourist Attractions for Desktop and Hand Held Devices," *Applied Artificial Intelligence*, vol. 17, no. 8-9, pp. 687-714, 2003.
- [29] S. Stoyanov, T. Glushkova, A. Stoyanova-Doycheva, J. Todorov and A. Toskova, "A Generic Architecture for Cyber-Physical-Social Space Applications," in *Intelligent Systems: Theory, Research and Innovation in Applications. Studies in Computational Intelligence*, vol. 864, Springer, Cham, 2020.
- [30] A. Sheth, P. Anantharam and C. Henson, "Physical-Cyber-Social Computing: An Early 21st Century Approach," *IEEE Intelligent Systems*, 2013.
- [31] U. Gretzel, M. Sigala, Z. Xiang and C. Koo, "Smart Tourism: Foundations and Developments," *Electronic Markets*, vol. 25, pp. 179-188, 2015.
- [32] F. Ali, L. Cain, T. S. Legendre and L. Wu, "The Intersection of Technology, Accessible Tourism and Tourists With Intellectual Disabilities: Proposing a Novel Conceptual Framework," *Journal of Hospitality & Tourism Research*, 2023.
- [33] J. Yoon and C. Choi, "Real-Time Context-Aware Recommendation System for Tourism," *Sensors*, 2023.
- [34] M. Florez, E. Carrillo, F. Mendes Neto and J. Carreno, "A Context-Aware Tourism Recommender System Using a Hybrid Method Combining Deep Learning and Ontology-Based Knowledge," *Journal of Theoretical and Applied Electronic Commerce Research*, 2025.
- [35] A. Smajić, R. Karlović, M. Bobanović Daško and I. Lorencin, "Context-Aware Tourism Recommendations Using Retrieval-Augmented Large Language Models and Semantic Re-Ranking," *Electronics*, vol. 14, no. 22, p. 4448, 2025.
- [36] P. Sahoo, P. Meharia, A. Ghosh, S. Saha, V. Jain and A. Chadha, "A Comprehensive Survey of Hallucination in Large Language, Image, Video and Audio Foundation Models," in *Findings of the Association for Computational Linguistics: EMNLP*, 2024.
- [37] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright and P. Mishkin, "Training Language Models to Follow Instructions with Human Feedback," in *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022.
- [38] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan and P. Dhariwal, "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 2020.
- [39] I. Nunes and D. Jannach, "A Systematic Review and Taxonomy of Explanations in Decision Support and Recommender Systems," *User Modeling and User-Adapted Interaction*, vol. 27, no. 3-5, pp. 393-444, 2017.
- [40] L. Atzori, A. Iera and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787-2805, 2010.
- [41] C. Perera, A. Zaslavsky, P. Christen and D. Georgakopoulos, "Context Aware Computing for the Internet of Things: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, pp. 414-454, 2014.
- [42] S. Stoyanov, A. Stoyanova-Doycheva, T. Glushkova and E. Doychev, "Virtual Physical Space: An Architecture Supporting Internet of Things Applications," in *20th International Symposium on Electrical Apparatus and Technologies (SIELA)*, 2018.
- [43] S. Stoyanov, A. Stoyanova-Doycheva, T. Glushkova, E. Doychev and J. Todorov, "A Reference Architecture of Internet of Things Ecosystem," *Computer Sciences and Communications*, 2018.
- [44] S. Stoyanov, T. Glushkova, E. Doychev, A. Stoyanova-Doycheva and V. Ivanova, *Cyber-Physical-Social Systems and Applications - Part 1*, Plovdiv: Plovdiv University Press, 2019.
- [45] T. Glushkova, S. Stoyanov, A. Stoyanova-Doycheva, V. Ivanova and L. Doukovska, "AmbiNet - An Environment for Ambient-Oriented Modeling," *International Journal of Computing*, 2019.
- [46] Z. Guglev and S. Stoyanov, "Hybrid Approach for Manipulation of Events in the Virtual Referent Space," in *International Scientific Conference "Blue Economy and Blue Development"*, 2018.
- [47] R. C. Martin, *Clean Architecture: A Craftsman's Guide to Software Structure and Design*, Prentice Hall, 2017.
- [48] E. Evans, *Domain-Driven Design: Tackling Complexity in the Heart of Software*, Addison-Wesley Professional, 2003.
- [49] L. Bass, P. Clements and R. Kazman, *Software Architecture in Practice*. 3rd ed, Addison-Wesley, 2012.
- [50] R. T. Fielding, "Architectural Styles and the Design of Network-based Software Architectures," PhD dissertation, University of California, Irvine, Irvine, CA, 2000.
- [51] M. T. Nygard, *Release It! Design and Deploy Production-Ready Software*, 2nd ed., Pragmatic Bookshelf, 2018.
- [52] M. Fowler, *Patterns of Enterprise Application Architecture*, Addison-Wesley Professional, 2002.

- [53] K. N. Milev, "Intelligent Personal Tourist Guide: dissertation for the award of the educational and scientific degree 'Doctor'," Plovdiv University "Paisii Hilendarski", FMI, COBISS.BG-ID 66731016, <https://plus.cobiss.net/cobiss/bg/bg/data/cobib/66731016>, Plovdiv, 2024.
- [54] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain Union legislative acts (Artificial Intelligence Act)," Official Journal of the European Union, L series, 12 July 2024, 2024.
- [55] G. Shani and A. Gunawardana, "Evaluating Recommendation Systems," in *Ricci, F., Rokach, L., Shapira, B., Kantor, P. B. (eds) Recommender Systems Handbook*, Springer, 2011.
- [56] Trips Journal, "Sights in Bulgaria," 2026. [Online]. Available: <https://tripsjournal.com/zabelejitelnosti>. [Accessed 15 April 2026].