Extending ISO-Space to Annotate Spatial Information in Route Descriptions

Samira Goudarzi Ali Akbari[1]
The University of Melbourne, Parkville, Australia

Stephan Winter
The University of Melbourne, Parkville, Australia

Martin Tomko
The University of Melbourne, Parkville, Australia

Abstract

A robust spatial annotation scheme that explicitly represents spatial elements and the links between them in English natural (human) language route descriptions is required to support computational interpretation. Yet, existing spatial annotation schemes for English do not capture all the spatial elements and links that occur in route descriptions, including perceptual expressions and reorientation actions. Using Formal Concept Analysis (FCA), the elements and links are derived by modeling the conceptual space of route descriptions. These elements and links are then compared against ISO-Space, an annotation standard for encoding spatial entities, motion events, and spatial relations [28], selected as the reference framework due to its relatively comprehensive coverage among existing schemes. Based on the analysis, targeted scheme refinements and extensions interoperable with ISO-Space are introduced. The resulting scheme is evaluated through two annotation rounds on 48 route descriptions drawn from four datasets. The results show improved consistency and increased annotation coverage through perceptual and reorientation links, which account for 23% of all annotated links, without compromising annotation reliability.

2012 ACM Subject Classification Computing methodologies → Language resources; Computing methodologies → Ontology engineering

Keywords and phrases spatial language, spatial information annotation scheme, route description

Digital Object Identifier 10.4230/LIPIcs.COSIT.2026.12

1 Introduction

People use verbal descriptions to communicate spatial information, thus externalizing their mental representations of their spatial environment. A route description is defined as "*a verbal instruction to follow a particular route through an environment*" [37, p. 4]. In real-world scenarios, such as emergency calls, route descriptions may be the only information available to guide someone to a specific location. The computational interpretation of such instructions supports applications including the conversion of verbal descriptions into sketch-map representations [34, 6] and geocoding [21]. These applications rely on the automatic extraction of spatial information from route descriptions. Such extraction typically involves machine learning models [30, 15, 27], which in turn require an annotated corpus for training.

To train a model for the automatic extraction of spatial information in route descriptions, we first need an annotation scheme that specifies which spatial elements and links to annotate. Consider the following example of a human-generated route description: *Walk straight along Park Street, and turn left when you see the library on your right. Follow the path past the stone bridge, then turn left at the fountain.* According to Denis [8], landmarks (e.g., *library*)

[1] Corresponding author

17th International Conference on Spatial Information Theory (COSIT 2026).
Editors: Sabine Timpf, Gabriele Filomena, Armand Kapaj, Rui Zhu, Nicholas Giudice, and Ed Manley;
Article No. 12; pp. 12:1–12:23
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and action descriptions constitute two critical elements of route descriptions. Denis further identifies two main classes of actions: changing orientation (e.g., *turn left*) and proceeding (e.g., *walk straight along*). In addition to the two main actions, Denis identifies another class of action, positioning and inception (e.g., *turn left when you see the library on your right*), in which perception verbs (e.g., *see*) participate [2]. Perceptual expressions allow the moving individual to check their current location and orientation, often based on descriptions of the visual scene, whereas changing orientation actions introduce decision points at which orientation needs to be changed. Perceptual expressions and changing orientation actions must be captured in an annotated dataset of route descriptions.

Generating annotated datasets requires robust annotation schemes that define which spatial elements and links can be consistently recognized across corpora. Previous research has proposed annotation schemes aimed at identifying spatial information in language [20, 16, 28] based on broader research in spatial semantics [31, 13]. None of these schemes fully captures the spatial elements that need to be annotated in route descriptions. For example, SpatialML [20] focuses primarily on the annotation of place elements but does not address the proceeding actions. Spatial Role Labeling [16] captures roles, but does not represent the full semantics of spatial relations expressed in spatial discourse. Pustejovsky et al. [28] go further by providing the ISO-Space annotation scheme to encode spatial entities, motion events, and spatial relations, aiming to ameliorate the limitations of earlier standards. However, critical elements frequently occurring in route descriptions, including perceptual expressions and reorientation actions, are still not explicitly represented through dedicated tags in the current ISO-Space specification. These gaps highlight the need for an annotation scheme that captures the spatial elements and links present in verbal route descriptions.

In this paper, three research questions (RQs) are addressed: **RQ1**: What conceptual elements and links are required to achieve a systematic representation of spatial information in walking route descriptions beyond what ISO-Space currently provides? **RQ2**: How can the unrepresented elements and links be formally defined and integrated into the ISO-Space framework without compromising its ontological coherence and interoperability? **RQ3**: How does the extended annotation scheme improve coverage and consistency compared with the original ISO-Space specification when applied to English route corpora?

To address RQ1, a systematic methodology based on Formal Concept Analysis (FCA) is used to model the conceptual space of spatial discourse and identify conceptual gaps in ISO-Space through a theory-driven lattice structure. This step demonstrates not only the conceptual gaps in ISO-Space scheme but also the need for refinements. For RQ2, each uncovered and refined spatial concept is formally defined and integrated into the ISO-Space framework using the CASCADES stages [7], ensuring ontological coherence and interoperability. Finally, RQ3 evaluates the extended scheme through two annotation rounds applied to three walking route corpora and one hiking route corpus: the first annotation round using the original ISO-Space standard and the second annotation round using the systematically extended standard, to assess improvements in coverage and consistency. These corpora represent naturalistic route descriptions, providing a suitable testbed to evaluate annotation coverage, and consistency. All datasets are in English, which limits the scope of this study to avoid cross-linguistic variation in spatial relations and grammatical patterns. The contributions of this study are therefore:

- Systematic identification of spatial elements and links in route descriptions that are not represented in the ISO-Space annotation standard, including reorientation actions and perception expressions;
- Systematic identification of spatial elements and links in route descriptions that require refinement in the ISO-Space annotation standard to improve coverage and consistency;

- Design of targeted extensions and refinements that explicitly encode missing elements and links, together with a formal XML specification that integrates these into a scheme compatible with the ISO-Space framework; and
- Application of the extended scheme to three corpora of walking route descriptions and one hiking corpus, demonstrating increased annotation coverage and consistency.

Section 2 reviews the conceptual foundations of route descriptions and existing spatial annotation schemes. Section 3 describes the datasets used to construct the annotation corpus. Section 4 applies Formal Concept Analysis to address the first research question. Section 5 details the integration of the resulting refinements and extensions into the ISO-Space XML representation, addressing the second research question. Section 6 describes the evaluation procedure and reports the results addressing the third research question. Section 7 concludes the paper and discusses implications, limitations, and future work.

2 Related work

This section provides an overview of the relevant literature. Section 2.1 examines the conceptual components of route descriptions and Section 2.2 reviews existing spatial annotation schemes.

2.1 Route description components

Route directions constitute a particular form of spatial discourse whose purpose is to guide an individual through an unfamiliar environment to a specified destination [9]. Other studies have identified the components involved in route descriptions and have proposed definitions for these components:

Segment. According to Tversky et al. [33], a segment is a conceptual unit that divides route depictions and descriptions into start point, reorientation, path/progression, and end point, which contains sufficient information to allow movement between points. Jackendoff's path constituents explicitly include start and end points of motion and decompose motion as a moving entity that traverses a path over a specific time period [12]. Allen [2] identifies pathways as environmental features that denote actual or potential channels of movement.

Decision point. Locations that offer alternative paths are typically exemplified by intersections [2]. In cities, a node is a strategic point or focus that a moving individual can enter and from which movement within the city is organized [19]. Nodes are junctions, which function as navigation decision points, or concentrations, which derive importance from the condensation of a specific use or physical character.

Landmark. Landmarks are point-references external to the moving entity, which does not enter them [19]. Based on Lynch, Allen [2] defines landmarks as environmental features that serve as point-references and intermediate goals along a path of movement from the origin to the destination. Denis [9] further characterizes landmarks as three-dimensional physical objects, either natural (e.g., *church*) or objects intended to convey directions (e.g., *signposts*), as well as two-dimensional entities (e.g., *streets*).

Action description. Denis [8] distinguishes two main action classes: proceeding (e.g., *walk straight ahead*) and changing orientation (e.g., *turn left*). Additional actions include positioning and inspection, which allow moving entities to check orientation based on the visual scene (e.g., *when you arrive there you will see the library on your right*) [8] and involve perception verbs (e.g., *see*) [2]. the proceeding and changing orientation are expressed by movement verbs, defined by Allen [2] as statements prescribing where the moving entity should go and semantically reducible to *go* or *turn*.

Unlike Allen, Winter et al. [37] distinguish verbs denoting proceeding actions from those expressing changes in orientation, the latter of which are not integrated into motion verb classes proposed in other literature. Muller [22] introduces six motion verb classes, including *leave, hit, reach, external, internal, and cross* within a qualitative theory of motion based on spatio-temporal primitives. Asher et al. [3] propose ten classes of motion verbs in French, including *s'approcher* (approach), *arriver* (arrive), *entrer* (enter), *se poser* (alight), *s'éloigner* (move away), *partir* (leave), *sortir* (exit), *décoller* (take off), *passer (par)* (pass through), and *dévier* (deviate). Pustejovsky et al. [29] map the motion classes proposed by Muller, Asher et al., and Baker et al. [5], resulting in ten unified motion verb classes: *move, move-external, move-internal, leave, reach, detach, hit, follow, deviate, and stay.*

Descriptive component. Spatial relations between environmental features or between moving entities and environmental features can be provided in route descriptions through direction designations [2] or topological relations [8]. This form of spatial description is conveyed using state-of-being verbs that are reducible to *is* [2]. Direction designations specify spatial relations using abstract, object-based, or body-based frames of reference [2].

Distance designation. Spaces that separate points of reference are expressed through standard units, conventional nonstandard units, temporal units, or vague judgments [2], and are expressed qualitatively rather than quantitatively [32].

This review of route description components establishes a theoretically grounded set of conceptual building blocks from prior research, which provides the basis to specify linguistic annotations units for evaluating the representational coverage of ISO-Space.

2.2 Spatial information annotation schemes

SpatialML is a framework developed to support the annotation of spatial information [20], mainly aimed at marking spatial locations for integration with external resources such as domain-specific databases and gazetteers. SpatialML does not represent motion events because it does not annotate predicates.

Kordjamshidi et al. [16] introduce Spatial Role Labeling (SpRL), a spatial annotation scheme grounded in Holistic Spatial Semantics (HSS) [38] that identifies spatial arguments and relations. SpRL models static and dynamic relations using three roles, including trajector, landmark, and spatial indicator, formally represented as semantic triples [16].

TimeML is created by Pustejovsky et al. to represent temporal information in text [26]. Pustejovsky et al. [29] discuss how the structure of spatial events motivates spatio-temporal markup, porposing SpatialML. ISO-Space, designed to capture spatio-temporal and spatial information in natural language, builds on the principles of TimeML and SpatialML [28]. Although ISO-Space aligns with SpRL, ISO-Space incorporates more detailed and fine-grained spatial and linguistic elements.

The detailed semantics of ISO-Space complicate data preparation for machine learning [15]. Yet, Pustejovsky et al. [28] argue for this detail. Simple schemes, such as SpRL may fail to capture the intended interpretation of spatial expressions. For example, assigning the topological relation Tangential Proper Part (TPP) to *The girl is in the room* does not clearly convey the intended meaning and may suggest an unintended interpretation, such as the girl being attached to the side of the room rather than standing on the floor inside it.

The latest version of ISO-Space (ISO 24617-7) provides an XML-based spatial information annotation scheme to represent spatial elements and links. The scheme defines elements for spatial entities, events, spatial relations, and dimensional information, the latter encoded using *<measure>*. Spatial entities comprise three subtypes: *<place>*, *<path>*, and *<eventPath>*. While *<place>* and *<path>* represent locational entities, *<eventPath>* denotes a trajectory created by a motion event. In addition to these locational or trajectorial entities, ISO-Space also allows for non-locational *<entity>* that function as a spatial entity or participates in motion events. Events are encoded using *<event>* and include both motion and non-motion events, with *<motion>* representing events that involve a change of location of an *<event>*. Spatial relations are encoded using *<sRelation>*, which includes attributes of *topological*, *directional*, *pathDefining*, and *goalDefining*. the topological and directional relations between elements are expressed using the link types *<qsLink>* and *<oLink>*, where *<oLink>* additionally incorporates an explicit frame of reference and reference point. Motion relations are encoded using *<moveLink>*, which relates a moving entity to the *<eventPath>* it traverses, while *<mLink>* associates a *<measure>* with the entity being measured.

ISO-Space (ISO 24617-7) is the only ISO standard for the annotation of spatial information. Both the standard and its accompanying corpus (SpaceEval) have been adopted in research on Natural Language Processing (NLP) to extract spatial relations [23, 30, 35, 24]. Despite this progress, it has not gained widespread adoption due to its complexity and the lack of annotated datasets and automated tagging tools [11]. This paper shows that this complexity does not hinder implementation and demonstrates that the proposed extensions and refinements improve coverage and consistency in annotating elements and links in route descriptions.

3 Data

This section details the data foundation and is divided into two subsections. Section 3.1 outlines the selection requirements to identify suitable datasets and introduces the selected datasets, describing their content and characteristics. Section 3.2 describes the procedure for compiling a representative corpus of 48 route descriptions.

3.1 Dataset selection and description

This research uses datasets of authentic route descriptions produced by different users, covering activity types (e.g., *walking, hiking*), geographic contexts (e.g., *urban, natural and rural*), and environmental settings (e.g., *indoor, outdoor*) to ensure generalization. Four datasets that meet these criteria are identified: *Walkingmaps*, *Slowways*, *Soleway*, and *Summitpost*. The first three describe walking activities, while *Summitpost* includes hiking and mountaineering activities. *Walkingmaps*, *Slowways*, and *Summitpost* represent outdoor route descriptions in rural, urban, and natural areas, whereas *Soleway* contains indoor ones.

- *Walkingmaps* (`https://walkingmaps.com.au/`) is a publicly accessible platform operated by Victoria Walks Inc., a non-profit organization that promotes walkable communities in Australia. It enables users to discover and share leisure walks. Each walk includes a trajectory, points of interest with verbal descriptions and optional images, and a textual description that provides navigational guidance and experience, along with metadata such as distance, estimated duration, and difficulty. Web scraping was used to extract 2,631 walks from the *Walkingmaps* website. The dataset was collected in July 2025.

Table 1 Size of the sampling frame across all datasets.

	Walkingmaps	Slowways	Soleway	Summitpost	Total
Count	2495	16229	1156	1795	21675

- The *Slowways* project (`https://beta.slowways.org/`) is a citizen-created national walking network connecting Britain's towns, cities, and national landscapes. Each route page includes an *Overview*, *Surveys*, *Geography*, and *Reviews*, with the option to download route geometry. The *Reviews* section provides textual descriptions with navigational information, while the *Overview* lists start and end point names, photos, and details. The *Surveys* section reports facilities, accessibility, obstacles, challenges, and key measurements, and the *Geography* section provides quantitative details such as length and elevation. Access to the dataset was granted in August 2025, comprising 16,450 routes.
- *Soleway* (`https://soleway.ugent.be/`) is a crowd-sourced indoor navigation service that offers wayfinding capabilities through textual route descriptions, avoiding expensive technologies, and promoting ease of use. Each route description includes metadata such as author name, author id, and created date. The dataset used in this study corresponds to the state of the *Soleway* platform in 2018 and contains 1,164 route descriptions.
- *Summitpost* (`https://www.summitpost.org/`) is a collaborative platform based in the united states that focuses on outdoor activities such as climbing. The search parameters for route entries include location (e.g., *continent*), route type (e.g., *hiking*), and rock difficulty. Route entries comprise sections such as *Overview*, *Getting There*, *Route Description*, *Essential Gear*, *Commentary*, *When to Climb*, and *Images*. Each route includes metadata such as coordinates, season of use, estimated duration, and difficulty ratings. In October 2025, 1,795 hiking route records met the search constraints of continent (North and South America), country (United States), and route type (hiking); these constraints were applied to ensure diversity of geographic contexts and activity types.

3.2 Corpus construction for annotation

The sampling frame consists of all available route descriptions from the four datasets. After excluding records that do not contain route descriptions, the final sampling frame comprises 21,675 descriptions, as summarized in Table 1. Systematic stratified sampling is used to capture the variability between heterogeneous strata of route descriptions from different datasets. The sampling frame is partitioned into mutually exclusive *strata*, and samples are drawn from each stratum so that all subgroups are represented. The stratified sampling procedure by which the final corpus of 48 route descriptions is derived from the sampling frame is described below.

- Each description is categorized by two criteria: geographic setting (urban, natural/rural, or indoor) and activity type (walking or hiking). The sampling frame is divided into strata defined by the cross-product of these criteria, yielding six strata. The indoor–hiking stratum is excluded as conceptually meaningless, while the urban–hiking stratum is removed due to sparse representation. The remaining strata are internally homogeneous with respect to geographic setting and activity type and together partition the sampling pool, ensuring valid combinations. The number of route descriptions in each stratum is summarized in Table 2.
- A quota-based sampling strategy is adopted, with a fixed quota of 12 route descriptions selected from each stratum, resulting in a corpus of 48 route descriptions (711 sentences). This corpus size supports reliable evaluation metrics for the training of computational

Table 2 Number of route descriptions in each stratum.

Stratum	Walkingmaps	Slowways	Soleway	Summitpost	Total
urban–walking	1656	0	0	0	1656
natural/rural–walking	839	16229	0	0	17068
indoor–walking	0	0	1156	0	1156
natural/rural–hiking	0	0	0	1795	1795

Table 3 Summary statistics for words and sentences in the 12 samples of the four strata.

	Words				Sentences			
Stratum	Count	Min	Mean	Max	Count	Min	Mean	Max
urban–walking	3396	114	283	486	260	12	22	39
natural/rural–walking	3032	117	253	472	209	10	17	42
indoor–walking	546	26	46	91	68	2	6	13
natural/rural–hiking	2432	70	203	361	174	4	15	26

models [30] and is consistent with the corpus sizes used to develop gold standard datasets [27, 1]. In proportionate allocation, the sample size per stratum is aligned with its relative size to preserve the population distribution and minimize bias. However, since more than 78.7% of the population falls within a single stratum, this approach is not suitable to achieve an interpretable sample. Optimum allocation requires prior knowledge of the variance per stratum, which, for textual data, depends on previously annotated elements that are unavailable. Therefore, minimum quotas are selected to ensure balanced representation across strata and avoid over-representing the dominant stratum.

- Within each stratum, random sampling is applied to select 12 route descriptions objectively and without bias. No contributor appears more than once in the final corpus; any violating description is replaced by another random one drawn from the same stratum. Table 3 presents summary statistics for the 12 route descriptions selected from each of the four strata. Count reports the total number of words or sentences per stratum, while Min, Mean, and Max summarize the per-description word and sentence counts. Indoor walking descriptions tend to be shorter because they focus on wayfinding within confined spaces, while outdoor walking, often undertaken as a leisure activity, typically takes hours and, consequently, involves longer route descriptions.

4 Scheme development via Formal Concept Analysis

To address RQ1, a systematic approach is required to explore the conceptual space of route descriptions. Formal Concept Analysis (FCA) is a mathematical theory of concept formation derived from lattice and ordered set theories that reveals relationships between objects and attributes [36]. FCA operates on a formal context, defined as a set of incidence relations between a set of objects and a set of attributes, indicating which objects possess which attributes. Its main construct is the *formal concept*, defined by its *extent* (the set of objects) and *intent* (the set of shared attributes). Formal concepts are organized through a generalization–specialization ordering that forms a concept lattice [36, 10]. Concept lattices can be represented as Hasse diagrams, which provide a structured visualization of conceptual spaces and support the systematic identification of conceptual gaps [10].

The process of applying FCA to answer RQ1 begins with a conceptual analysis in Section 4.1, which specifies the semantic requirements of the route descriptions. Building on this foundation, Section 4.2 describes the construction of formal contexts for both *elements* and *links*, followed by Section 4.3, which details the generation of the corresponding concept lattices. The resulting structures are then used in Section 4.4 to propose refinements and in Section 4.5 to extend newly introduced elements and links that address the conceptual gaps identified through FCA.

4.1 Conceptual analysis

The linguistic units in the annotation scheme are either *elements*, representing semantic arguments in natural (human) language route descriptions, or *links*, expressing relations between those elements. Conceptual analysis serves as the primary analytical stage, specifying the linguistic units required to capture route description components. In the following, the linguistic units, including elements and links, are defined.

- **Elements**:
  place: A word(s) that refers to an entity resulting from a shared identification of a location. This definition is based on the derivation that reflects an ontological commitment introduced by Purves et al. [25]. This linguistic unit is used to annotate references to landmarks and decision points based on the notion of Allen [2].
  path: A word(s) that denotes a linear entity that functions as an actual or potential traversable channel. This linguistic unit is used to annotate references to paths following Allen [2].
  moving entity: A word(s) that denotes the entity undergoing a change of location.
  spatial relation: A word(s) that indicates how two entities are located relative to each other. In syntactic terms, spatial relations are most often represented using prepositions, but not always. This linguistic unit is used to annotate references to direction designations [2] or topological relations [8].
  motion: A verb that indicates an event in which the moving entity moves from a source, along a path, and towards a goal. This linguistic unit is used to annotate references to motion verbs [29] in proceeding actions [8].
  perception: A verb that indicates an event in which the moving entity observes or notices an entity, using sensory information to confirm location or orientation. This linguistic unit is used to annotate references to perception verbs [37] in inspection and positioning actions [8].
  reorientation: A verb that indicates an event in which the moving entity changes facing direction without displacement or with displacement, typically at decision points. This linguistic unit is used to annotate references to reorientation verbs [37] in changing orientation actions [8].
  event: A word(s) that indicates a broad notion of an event that includes all kinds of actions, states, processes, activities, and accomplishments [4]. This linguistic unit is used to annotate references to events, except perception, reorientation, and motion events.
  measure: A word(s) that indicates the degree of separation between two reference points qualitatively and quantitatively and applies to both spatial and temporal dimensions. This linguistic unit is used to annotate references to distance designations [2].
- **Roles**:
  figure: An entity whose location, orientation, or motion is described.
  ground: An entity relative to which the location, orientation, or motion of the figure is specified.
- **Links**:
  topological link: A link that associates a *figure* with a *ground* through a topological spatial relation. This linguistic unit is used to relate elements in descriptive components [2].

directional link: A link that associates a *figure* with a *ground* through a directional spatial relation within a frame of reference (absolute, intrinsic, relative). This linguistic unit is used to relate elements in descriptive components [2].
movement link: A link that associates a motion event with its *figure*, *ground*, and involved elements of *source* (path or place that is the origin of the movement), *path* (path in which movement occurs), *mids* (path or place that is intermediate goals of movement), *spatial relation*, and *measure*. This linguistic unit is used to relate the elements in proceeding actions [8].
reorient link: A link that associates a reorientation event with its *figure*, *ground*, and involved elements of *source* (the place or path at which the reorientation occurs), and *spatial relation*, indicating the new facing direction. This linguistic unit is used to relate the elements in changing orientation actions [8].
perceive link: A link that associates a perception event with its *figure* and *ground*, recording sensory confirmation of a *ground*. This linguistic unit is used to relate elements in inspection and positioning actions [8].

The defined units guide the determination of the sets of objects and attributes required to construct the formal contexts, with elements and links treated separately.

4.2 Formal context construction

The construction of formal contexts begins by determining the sets of objects and attributes required to represent the units defined in Section 4.1 separately for elements and links. The objects constitute the rows, and the attributes constitute the columns of the context matrices. Each cell is assigned the value 1 when an object exhibits a given attribute and 0 otherwise.

All meaningful combinations of attributes are systematically generated to support comparison across objects and to identify three outcomes: combinations missing from ISO-Space, combinations that correspond to existing ISO-Space elements and links, and combinations for which an ISO-Space counterpart exists but requires refinement. Semantically meaningless combinations are excluded. For missing combinations, a new object name is introduced; for corresponding combinations, the ISO-Space name is adopted; and for incomplete combinations, the ISO-Space name is retained, and the object is treated as a refinement.

To ensure a consistent and readable presentation, a uniform notation is adopted for all objects and attributes in the FCA. The names of objects, whether *elements* or *links*, are noted using angle brackets (e.g., *<entity>*, *<qsLink>*) and attributes are prefixed with @ (e.g., *@event*, *@spatialrelation*), following the convention used in the ISO standard. To distinguish the contributions of this work from the original ISO-Space specification, all newly introduced or modified attributes and objects are marked with the *rs:* prefix (***R****oute* ***S****emantics*), following standard XML conventions. For example, while *<qslink>* remains an ISO-Space link, *<rs:reorientLink>* represents a new link introduced in this work. In the following, the construction of the formal contexts for both elements and links is described. The modified ISO-Space scheme, incorporating both the proposed extensions and refinements, is referred to as *ISO-Space+RP*, where *R* denotes reorientation and *P* denotes perception.

- **Elements:** A set of attributes is selected to conceptually characterize the semantics of objects. Attributes already present in ISO-Space, such as *@spatial entity*, *@locational-one*, *@locational-series*, *@non-locational*, *@traversal*, *@event*, *@motion event*, *@motionClass*, *@non-motion event*, *@spatial relation*, *@topological*, *@directional*, *@pathDefining*, and *@goalDefining*, are reused whenever they correspond to a required semantic distinction. Two further distinctions of events required lack ISO-Space counterparts. Therefore, we

Table 4 Formal context for elements.

	@spatial entity	*@locational-one*	*@locational-series*	*@non locational*	*@traversal*	*@event*	*@motion event*	*@motionClass*	*@non motion event*	*@rs:reorientation event*	*@rs:perception event*	*@spatial relation*	*@topological*	*@directional*	*@pathDefining*	*@goalDefinig*
<entity>	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
<rs:place>	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
<rs:path>	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0
<eventPath>	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0
<motion>	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0
<event>	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0
<sRelation>	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
<rs:reorientation>	0	0	0	0	0	1	0	0	1	1	0	0	0	0	0	0
<rs:perception>	0	0	0	0	0	1	0	0	1	0	1	0	0	0	0	0

introduce two additional attributes *@rs:reorientation event* and *@rs:perception event* to ensure that all meaningful attribute combinations required to represent elements in the route descriptions can be captured systematically.

The objects *<entity>*, *<rs:place>*, *<rs:path>*, *<eventPath>*, *<motion>*, *<event>*, and *<sRelation>* of ISO-Space are semantically aligned or require modification to align with the required elements defined in Section 4.1. The meaningful attribute combinations and conceptual analysis show that route descriptions introduce additional events, such as *<rs:reorientation>* (e.g., *turn* in *turn left at the fountain*) and *<rs:perception>* (e.g., *see* in *you see the library on your right*), for which ISO-Space provides no specific elements. Table 4 shows the formal context of the elements.

- **Links:** For attributes, whenever a required distinction already exists in ISO-Space, its attribute name is reused; these include *@event*, *@spatialRelation*, *@figure*, *@ground*, *@trelType*, *@drelType*, *@frameType*, *@mrelType*, *@motion*, *@sRelation*, *@merelType*, *@measure*, and *@bound*. When conceptual analysis reveals distinctions that lack ISO-Space counterparts, a new attribute name is introduced. Consequently, two additional attributes, *@rs:reorientation* and *@rs:perception*, are added to the attribute set to enable the construction of *<rs:reorientLink>* and *<rs:perceiveLink>*.

  For objects, combinations of attributes that fully correspond to existing ISO-Space links adopt the ISO-Space name, as in *<qsLink>*, *<oLink>*, and *<mLink>*. Combinations that partially correspond retain the ISO-Space name but are treated as refinements, as illustrated by *<rs:moveLink>*. Combinations without an ISO-Space counterpart require new links, motivating the introduction of *<rs:reorientLink>* and *<rs:perceiveLink>*. Table 5 presents the formal context for links.

The formal contexts resulting from both elements and links shown in Tables 4 and 5 provide a systematic representation of how objects are distributed across semantic attributes, establishing a common analytical basis for constructing the concept lattices. The lattices enable the enumeration of all logically valid attribute combinations, the evaluation of the ISO-Space coverage for elements and links, and the identification of required extensions or refinements to represent the full conceptual space of route descriptions in walking context.

Table 5 Formal context for links.

	@event	*@spatialrelation*	*@figure*	*@ground*	*@source*	*@mids*	*@path*	*@trelType*	*@drelType*	*@frameType*	*@mrelType*	*@motion*	*@rs:reorientation*	*@sRelation*	*@rs:perception*	*@mereltype*	*@measure*	*@bound*
<qsLink>	0	1	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0
<oLink>	0	1	1	1	0	0	0	0	1	1	0	0	0	0	0	0	0	0
<mLink>	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1
<rs:moveLink>	1	0	1	1	1	1	1	0	0	0	1	1	0	1	0	0	1	0
<rs:reorientLink>	1	0	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0
<rs:perceiveLink>	1	0	1	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0

4.3 Concept lattice construction

The formal contexts constructed in the previous subsection provide the input for computing the formal concepts for elements and links following standard definitions. Each concept is obtained by computing the closure operation of a set of objects or attributes, resulting in a pair consisting of an extent and an intent. Ordering these formal concepts by extent inclusion produces a complete concept lattice. Each node corresponds to a pairing of an extent and an intent, reflecting how elements cluster according to shared attributes. The concept lattices for elements and links are visualized using Hasse diagrams in Figures 1 and 2.

In the Hasse diagrams (Figures 1 and 2), each node represents a formal concept, with the extent shown in red below the node and the intent in dark blue above it. Concepts absent from ISO-Space are highlighted in purple, those requiring refinement in light yellow, and all others remain uncolored.

- **Elements:** For the *element* formal context, all closed object–attribute pairs were computed and ordered by extent inclusion to obtain the concept lattice shown in Figure 1. This lattice illustrates how elements inherited from ISO-Space, along with newly introduced ones, are organized into a hierarchy based on shared attributes.
- **Links:** For *links*, the formal context is processed analogously: all formal concepts are computed and ordered to form the concept lattice shown in Figure 2. This lattice reveals how existing ISO-Space links and newly proposed links are positioned within the conceptual space based on their attributes.

4.4 Refinements

In the following, the refinements for elements and links are completely described.

- **Elements:** According to Section 3.15 of ISO 24617-7, a *<path>* is defined as *“a location where the focus is on the potential for traversal or which functions as a boundary”*. The specification further notes that *“a <path> might be represented as an undirected graph whose vertices are locations and whose edges signify continuity; i.e., unlike an <eventPath>, a path has no inherent directionality”*. Section 3.3 defines *<eventPath>* as a *“directed path followed by a mover and coincident with a motion-event”*, noting that *“unlike (static) paths such as roads or circular tracks, event-paths are each triggered by a specific motion-event, characterized as being finite directed paths each with a start and an end”*, and further specifies that an *<eventPath>* is a non-consuming tag. A non-consuming tag does not anchor to textual spans, so its *@target* attribute is null. *@target* is an attribute of *<eventPath>* that maintains the textual span.

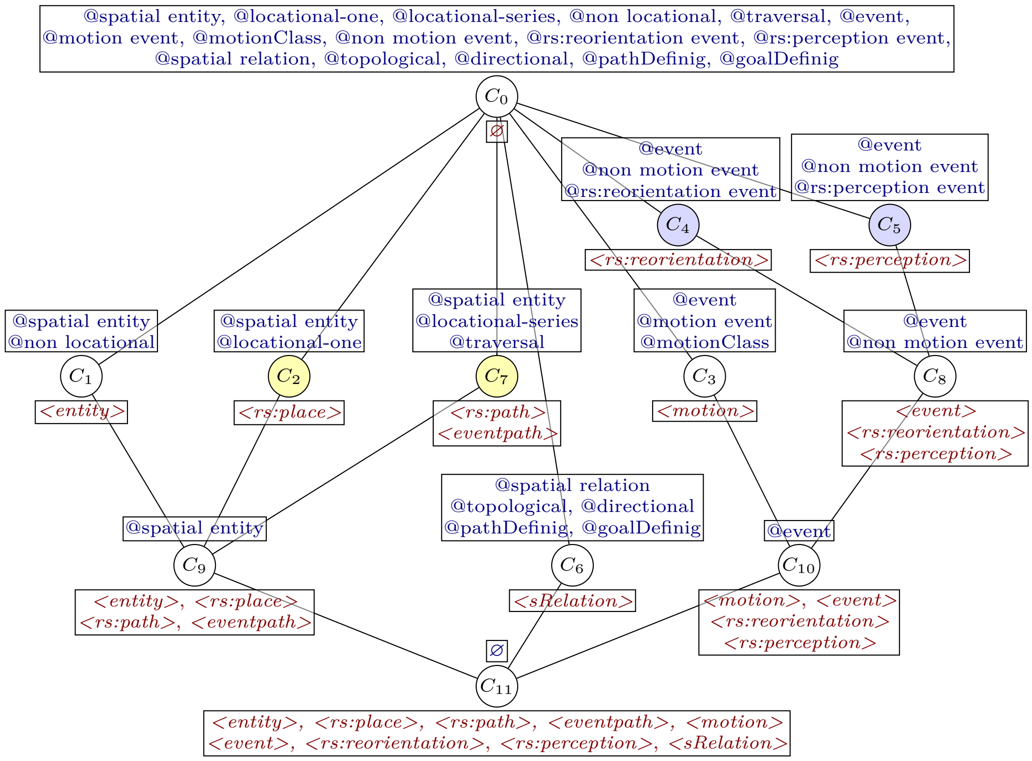

Figure 1 Concept lattice for elements: each node represents a formal concept with an extent (elements below, dark red) and an intent (attributes above, dark blue), ordered by extent inclusion. Nodes absent from ISO-Space are shown in purple, and nodes requiring refinement in yellow.

<eventPath> assumes that the path of movement is implicitly recoverable; however, in verbal route descriptions it is often expressed explicitly (e.g., *walk straight along Park Street*). This raises an annotation problem: should an explicit path of movement be annotated as *<path>* or *<eventPath>*? Such paths fit neither the non-consuming definition of *<eventPath>* nor the notion of an undirected *<path>*.

To resolve this disambiguation, either *<eventPath>* must not be non-consuming, or *<path>* must be directional. Representing movement paths via a non-consuming *<eventPath>* complicates annotation and downstream machine-learning tasks; therefore, *<eventPath>* is eliminated and *<path>* is redefined to include direction. Consequently, *<rs:path>* and *<eventPath>* co-occur in node *C7* with identical attributes. Under the refined definition of *<rs:path>*, *<eventPath>* has no distinct attributes, indicating that the two elements are indistinguishable and should be merged.

The next refinement concerns *<place>*. In Section 3.16 of the ISO 24617-7, *<place>* is defined as *"a geographic or administrative entity that is situated at a location"*, annotating *"geographic entities such as lakes and mountains, as well as administrative entities such as towns and counties"*. Under this definition, not all words referring to entities with a shared identification of location can be annotated as *<place>* (e.g., *fountain* in *turn left at the fountain*), limiting its use for place references in route descriptions. In addition, the *@type* attribute of *<place>* is restricted to predefined values and does not allow the annotation of all words that align with the definition in Section 4.1. Therefore, the definition of *<place>* and its *@type* attribute are modified to allow *<rs:place>* to annotate place references in route descriptions.

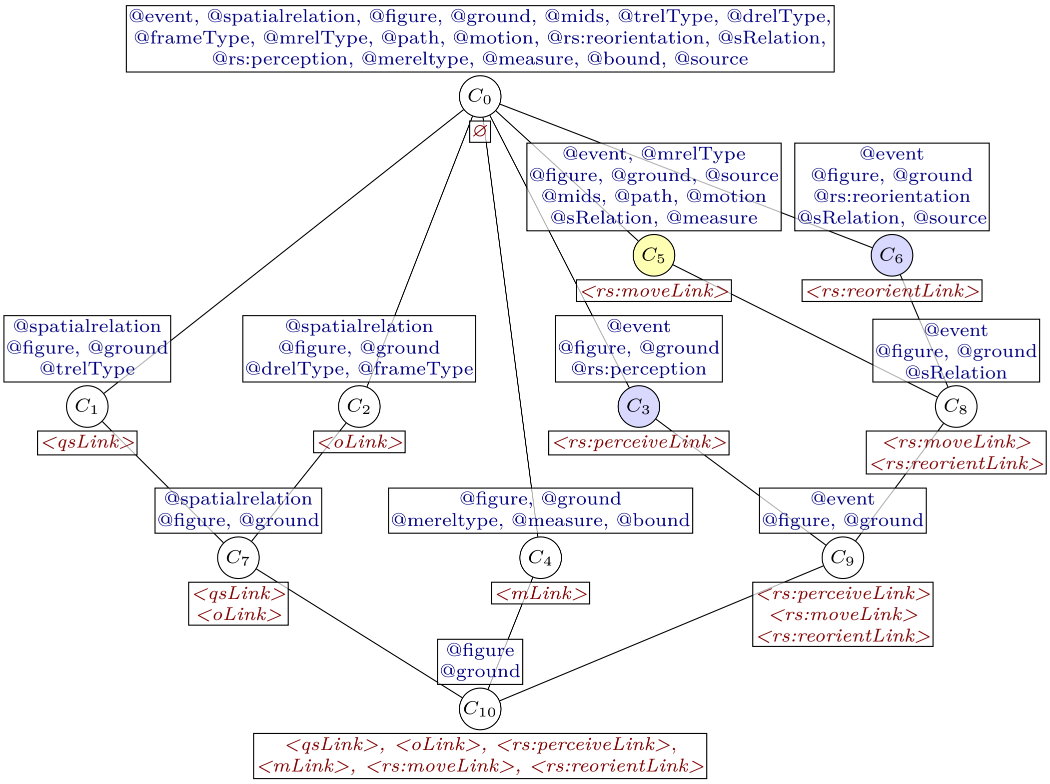

Figure 2 Concept lattice for links: each node represents a formal concept with an extent (links below, dark red) and an intent (attributes above, dark blue), ordered by extent inclusion. Nodes absent from ISO-Space are shown in purple, and nodes requiring refinement in yellow.

- **Links:** According to Section 8.4.3 of ISO 24617-7, *<moveLink>* is defined as a triple $\langle \eta, E, \rho \rangle$, where η is the mover and figure, E is a singleton containing a *<eventPath>* entity serving as the ground, and ρ is the relation between η and E triggered by a *<motion>* event. Its XML attributes include an identifier, a *@relType* (default *traverse* or a *@motionClass*), the figure, the ground, and the trigger. Refinement of *<moveLink>* is motivated by four issues: elimination of *<eventPath>* requires encoding the path of movement, the attributes lack directional information and *<measure>*, and they do not encode the source, mids, or goal of movement. Accordingly, node *C5* in the link concept lattice adds *@sRelation*, *@measure*, *@mids*, *@source*, and *@ground* (the goal of movement), which are absent from the current specification.

4.5 Extensions

Extensions are suggested for situations where concept lattices reveal combinations of attributes not captured by the original specification.

- **Elements:** According to the definition provided in Section 3.2 of the ISO 24617-7 standard, an *<event>* is defined as *"something that can be said to obtain or hold true, to happen, or to occur"*. In the abstract syntax, events constitute the set E, including both motions and non-motions. Section 8.3.6 further clarifies that a non-motion *<event>* corresponds to a TimeML *<event>* that does not involve a change of location.

Section 3.10 defines *<motion>* as an *"action or process involving the translocation of a spatial object, transformation of some spatial property of an object, or change in the conformation of an object"*. In the abstract syntax, motions belong to the set E but are distinguished from non-motions. Section 8.3.5 specifies that a *<motion>* is *"a TimeML event that involves a change of location of some entity"*. When an event is recognized as a motion, *<event>* is replaced by *<motion>* with attributes *@motionType*, *@motionManner*, *@motionGoal*, *@motionClass*, and *@motionSense*, supporting distinctions in manner vs. path, goals, and literal, fictive, or intrinsic interpretations.

Although route descriptions frequently contain reorientation verbs, the definition of *<motion>* excludes them. Reorientation verbs describe a change in facing direction and correspond to Denis's action class of changing orientation [8]. Section 3.10 restricts *<motion>* to events involving translocation, transformation of spatial property, or change in conformation, none of which covers a change in facing direction without displacement or with displacement. Consequently, reorientation verbs fall outside the scope of *<motion>* in ISO-Space and cannot be annotated as such without distorting their semantics. Under the generic definition of *<event>* in Section 3.2, these verbs would be annotated as *<event>*, obscuring their interpretive role; therefore, route descriptions require a distinct tag to capture their semantics.

Perceptual experiences function as descriptive components in route descriptions [2], introducing landmarks along the route or visible from a distance. They participate in the positioning and inspection actions by Denis [8], allowing a moving entity to check orientation based on the visual scene. Under the definition of *<event>* in Section 3.2, these verbs fall under *<event>* but require a dedicated tag distinct from *<event>*.

The concept lattice shows that neither *<reorientation>* nor *<perception>* shares all attributes with existing elements such as *<event>* and *<motion>*, appearing instead in singleton nodes (*C4* and *C5*). As they do not cluster with any ISO-Space elements, they constitute new concepts, leading to an extension of the element taxonomy with *<rs:reorientation>* and *<rs:perception>*. This extension ensures that route instructions involving turns or perceptual cues are explicitly represented and linked through *<rs:reorientLink>* and *<rs:perceiveLink>* to their figure and ground.

- **Links:** The lattice contains singleton concepts (*C3* and *C6*) whose extents consist only of links added to capture attribute combinations, namely *<rs:reorientLink>* and *<rs:perceiveLink>*. These appear as independent nodes because their attributes, *@rs:reorientation* and *@rs:perception*, do not co-occur with any links. Their position reflects that these attribute combinations are conceptually valid yet absent from ISO-Space; therefore, *<rs:reorientLink>* and *<rs:perceiveLink>* complete the conceptual space of links.

5 Integration of refinements and extensions via CASCADES

To answer RQ2, the CASCADES methodology (**C**onceptual analysis, **A**bstract syntax, **S**emantics, and **C**oncrete syntax for **A**nnotation **DES**ign) is used to integrate missing spatial concepts. CASCADES, proposed by Bunt [7], divides annotation design into four stages: 1) *Conceptual Analysis*, which identifies the information to be captured by analyzing relevant linguistic and cognitive theories; 2) *Abstract Syntax*, which defines the vocabulary of basic concepts and their possible combinations as set-theoretic constructs; 3) *Semantics*, which provides a formal interpretation of the abstract syntax; and 4) *Concrete Syntax*, which specifies how annotation structures are represented in a concrete format.

CASCADES provides a procedural framework for translating conceptual analysis into ISO-Space+RP specifications. The refinements and extensions identified in Section 4 are formalized in ISO-Space abstract syntax and XML-based concrete syntax as a machine-readable representation, with DTD declarations provided in Appendix A.

- **Refinements**
  <rs:path>: The attributes of this element remain unchanged from the original standard.
  <rs:place>: The attributes remain unchanged except for *@type*, which is treated as *CDATA* under the modified definition.
  <rs:moveLink>: The attributes include *@id*, *@reltype*, *@figure*, *@ground*, *@source*, *@mids*, *@path*, *@trigger*, *@sRelation*, and *@measure*. The refinements update *@ground* to refer to the goal of movement and add the attributes *@sRelation*, *@measure*, *@source*, *@mids*, and *@path*.
- **Extensions**
  <rs:reorientation> and **<rs:perception>**: The attributes of these elements include *@id* and *@target*.
  <rs:reorientLink>: The attributes include *@id*, *@figure*, *@ground* (the place or path used as a reference for the change in direction), *@source* (the place or path at which the reorientation occurs), *@trigger*, and *@sRelation*.
  <rs:perceiveLink>: The attributes include *@id*, *@figure*, *@ground*, and *@trigger*.

6 Evaluation of the modified annotation scheme

Addressing RQ3, this section presents a systematic evaluation of the ISO-Space+RP scheme for route descriptions. The evaluation assesses whether the ISO-Space+RP scheme improves coverage and consistency and is conducted on the corpus described in Section 3.2. Two annotation rounds are performed, each involving two annotators who are PhD students with research backgrounds relevant to geographic information science, using the INCEpTION annotation tool [14]. Two annotation guidelines are used: the first and second rounds follow ISO-Space and ISO-Space+RP specifications for elements, links, and attributes, respectively.

In the first annotation round, the corpus is annotated using the original ISO-Space specifications and inter-annotator agreement (IAA) is computed; in the second round, it is annotated using the ISO-Space+RP scheme and IAA recalculated. The results are compared using two metrics: *coverage*, defined as the number and proportion of identified elements and links, and *consistency*, measured by Krippendorff's α [17]. Agreement values between 0.41 and 0.60 are interpreted as moderate, values between 0.61 and 0.80 as substantial, and values above 0.81 as almost perfect agreement [18]. Element agreement is computed at the span level, requiring overlapping spans with the same element type, while link agreement applies a strict criterion requiring all link attributes to match or overlap between annotators. The evaluation results, comprising the count and proportion of elements and links (Figure 3) and inter-annotator agreement (Figure 4), with an emphasis on the proposed refinements and extensions, are discussed in the following.

In the first annotation, *<eventPath>* accounts for 11% of the annotated elements and is removed in the ISO-Space+RP scheme without loss of spatial information, as it does not anchor to explicit token spans. Its attributes, *@mids*, *@source*, and *@ground*, are instead represented in *<rs:moveLink>*, ensuring no loss of spatial information. The agreement at the element level is lowest for *<eventPath>* ($\alpha = 0.28$), indicating substantial ambiguity between the two *<eventPath>* and *<path>* discussed in Section 4.4. This ambiguity is further reflected in the counts: although *<eventPath>* is required as the *@ground* of *<moveLink>* in the original scheme, the corpus contains 623 *<moveLink>* but only 403 *<eventPath>*.

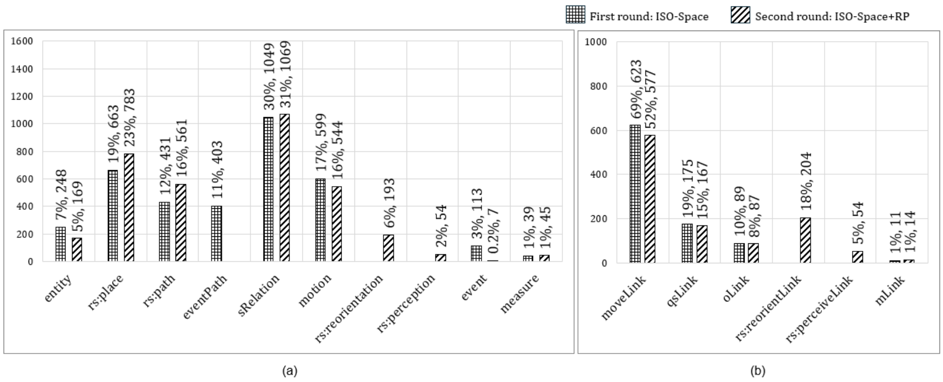

Figure 3 Statistics for two annotation rounds: (a) elements and (b) links. Percentages and counts are shown above bars.

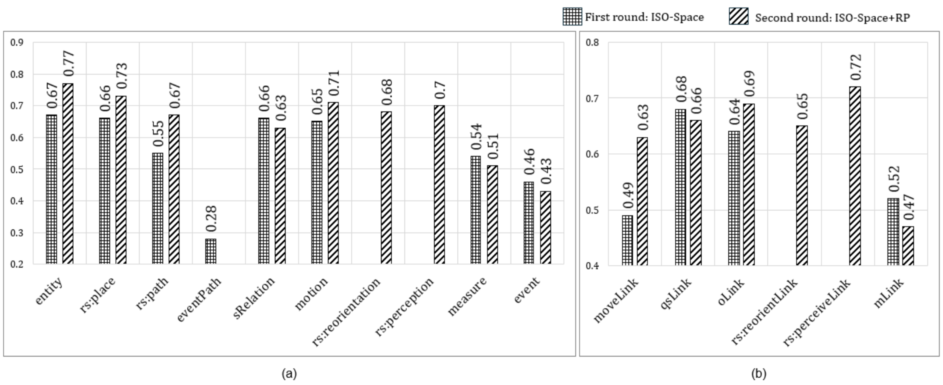

Figure 4 IAA for two annotation rounds: (a) elements and (b) links. Values are shown above each bar.

The number of *<rs:path>* increases from 431 to 561 due to the elimination of *<eventPath>* and the revised definition of *<rs:path>*. In the first annotation, the definition of *<path>* does not allow annotation of explicitly mentioned movement paths, leading to under-annotation, and such references cannot be annotated using the non-consuming *<eventPath>* tag. Under the revised definition, movement paths are annotated as *<rs:path>* in the second annotation. The revised definition and elimination of *<eventPath>* increase the IAA of *<rs:path>* from 0.55 to 0.67 by removing ambiguity in annotating movement path references.

The number of *<rs:place>* annotations increases from 663 to 783. Due to the revised definition of *<rs:place>*, all words that refer to entities with a shared identification of location are annotated as *<rs:place>*. Consequently, *<rs:place>* IAA increases from 0.66 to 0.73.

In ISO-Space scheme, *<moveLink>* includes three attributes, *@figure*, *@motion*, and *@ground*, with *<eventPath>* as the permissible *@ground*. The following statistics demonstrate that modified *<rs:moveLink>* is critical: 29 of 54 *<measure>* (64%), 368 of 561 *<path>* (66%), 556 of 1,069 *<sRelation>* (52%), and 370 of 783 *<place>* (48%) occur within *<rs:moveLink>* but are ignored by ISO-Space scheme. Although *@source*, *@mids*, and *@ground* are attributes of *<eventPath>* in ISO-Space scheme, the high inconsistency of *<eventPath>* ($\alpha = 0.28$) undermines its suitability to hold these attributes. At the link level, *<moveLink>* shows the lowest agreement ($\alpha = 0.49$) in the first round, increasing to $\alpha = 0.63$ in the second round.

In the second annotation, introducing *<rs:perception>* and *<rs:reorientation>* enables 247 elements (7%) and 258 corresponding links, *<rs:perceiveLink>* and *<rs:reorientLink>*, accounting for 23% of all links. Both *<perception>* and *<reorientation>* occur in all four datasets. Across the 12 route descriptions per stratum (urban-walking, natural/rural-walking, indoor-walking, and natural/rural-hiking), counts of *<perception>* are 30, 11, 5, and 54, and counts of *<reorientation>* are 59, 66, 32, and 36, respectively. These results show that the extensions are necessary in walking contexts, as these elements occur consistently in route descriptions regardless of activity type or environmental context. The introduced elements *<rs:perception>* ($\alpha = 0.70$) and *<rs:reorientation>* ($\alpha = 0.68$), and their corresponding links *<rs:perceiveLink>* ($\alpha = 0.72$) and *<rs:reorientLink>* ($\alpha = 0.65$), show substantial agreement.

The total number of elements increases from 3,142 to 3,425, and links from 898 to 1,103. The increases in elements are due to *<rs:reorientation>*, *<rs:perception>*, *<rs:path>*, and *<rs:place>*, while the counts for *<event>*, *<motion>*, and *<entity>* decrease. This decrease does not indicate information loss; instances previously annotated as *<event>* (104 of 113) and *<motion>* (63 of 599) are reclassified as *<rs:reorientation>*, and instances previously annotated as *<entity>* (106 of 248) are reclassified as *<rs:place>*. The increase in link counts results from the introduction of *<rs:perceiveLink>* and *<rs:reorientLink>*, while the number of *<rs:moveLink>* decreases as 63 of 623 instances are reassigned to *<rs:reorientLink>*.

Overall IAA increases from 58% to 67% for elements and from 55% to 66% for links. These gains occur despite reduced agreement for some elements, including *<sRelation>*, *<event>*, and *<measure>*, and for the links *<qsLink>* and *<mLink>*. Refinement of *<rs:place>* and *<rs:path>* and elimination of *<eventPath>* contribute to higher agreement, while restructuring *<rs:moveLink>* also increases IAA despite requiring agreement on more attributes. New introduced elements and links show substantial IAA.

7 Conclusion

Addressing research questions. This paper identifies the spatial elements and links required to represent route descriptions in walikng and hiking contexts. Using Formal Concept Analysis, the study identifies conceptual gaps in ISO-Space relevant to route descriptions, including perceptual expressions and reorientation actions, thus addressing RQ1. These concepts are defined and integrated into the ISO-Space framework through targeted extensions and refinements, using a compatible XML representation that preserves ontological coherence and interoperability, thus addressing RQ2. Two annotation rounds evaluate the impact of the targeted extensions and refinements on inter-annotator agreement and annotation coverage. The results show that ISO-Space+RP scheme increases the annotation coverage without compromising reliability, thus addressing RQ3.

Interpretation and contributions. The increases in element and link counts result from the refinements and extensions and reflect reclassification rather than information loss. The introduction of extensions enables the annotation of *<rs:perception>* and *<rs:reorientation>* and their associated links that occur in route descriptions regardless of the type of activity including walking and hiking or environmental context including urban, natural, rural and indoor. Agreement increases for both elements and links, demonstrating that the refinements improve consistency and that the extensions maintain substantial inter-annotator agreement while remaining compatible with ISO-Space. The ISO-Space+RP scheme enables the training of machine learning models that automatically detect spatial elements and links in walking and hiking contexts beyond those identifiable under ISO-Space. Specifically, the *<rs:reorientation>* element enables explicit and computational segmentation of scenes, determined by the position and viewing orientation of a moving entity. Segmentation of scenes is required for the interpretation of directional spatial relations. The *<rs:perception>* element introduces visible places that support the interpretation of spatial relations.

Limitations and future work. Some elements and links, including *<measure>*, *<event>*, and *<mLink>*, remain sparse, potentially affecting the stability of the agreement. The evaluation dataset is also restricted to walking and hiking. Route descriptions from other contexts may reveal other conceptual elements and links beyond those proposed in this study. Future work could apply the ISO-Space+RP scheme to broader activity types, add adjudication to resolve disagreements and improve reliability, expand the corpus to better represent under-represented elements and links including *<measure>*, *<event>*, and *<mLink>*, and investigate the impact of ISO-Space+RP scheme on downstream tasks such as route reconstruction and spatial reasoning.

References

1 Niloofar Aflaki, Shaun Russell, and Kristin Stock. Challenges in creating an annotated set of geospatial natural language descriptions (short paper). In *10th International Conference on Geographic Information Science (GIScience 2018)*, pages 20:1–20:6. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.GISCIENCE.2018.20.

2 Gary L Allen. Principles and practices for communicating route knowledge. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 14(4):333–359, 2000.

3 Nicholas Asher and Pierre Sablayrolles. A typology and discourse semantics for motion verbs and spatial PPS in French. *Journal of Semantics*, 12(2):163–209, April 1995. doi:10.1093/jos/12.2.163.

4 Emmon Bach. The algebra of events. *Linguistics and Philosophy*, 9(1):5–16, 1986. doi:10.1007/BF00627432.

5 Collin F. Baker, Charles J. Fillmore, and John B. Lowe. The Berkeley FrameNet project. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada, August 1998. Association for Computational Linguistics. doi:10.3115/980845.980860.

6 Lamia Belouaer, David Brosset, and Christophe Claramunt. From verbal route descriptions to sketch maps in natural environments. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2996913.2997003.

7 Harry Bunt. A methodology for designing semantic annotations. Technical Report TR 2013-001, TiCC Technical Report, Tilburg Center for Cognition and Communication, Tilburg University, 2013. URL: https://let.uvt.nl/general/people/bunt/docs/CASCADES-TiCC-TR.pdf.

8 Michel Denis. The description of routes: A cognitive approach to the production of spatial discourse. *Current Psychology of Cognition*, 16:409–458, 1997.

9 Michel Denis. *Space and spatial cognition: A multidisciplinary perspective*. Routledge, 2017.

10 Bernhard Ganter, Gerd Stumme, and Rudolf Wille. Formal concept analysis: Theory and applications. *JUCS - Journal of Universal Computer Science*, 10(8):926–926, 2004. doi:10.3217/jucs-010-08.

11 Alexander Henlein, Giuseppe Abrami, Attila Kett, and Alexander Mehler. Transfer of ISO-Space into a 3D environment for annotations and applications. In Harry Bunt, editor, *Proceedings of the 16th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pages 32–35, Marseille, May 2020. European Language Resources Association. URL: https://aclanthology.org/2020.isa-1.4/.

12 Ray Jackendoff. Conceptual semantics and cognitive linguistics. *Cognitive Linguistics*, 7(1):93–129, 1996. doi:10.1515/cogl.1996.7.1.93.

13 Olga Kagan. *The semantics of case*. Key Topics in Semantics and Pragmatics. Cambridge University Press, Cambridge, 2020. doi:10.1017/9781108236867.

14 Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9. Association for Computational Linguistics, June 2018. The 27th International Conference on Computational Linguistics (COLING 2018). URL: http://tubiblio.ulb.tu-darmstadt.de/106270/.

15 Parisa Kordjamshidi, Steven Bethard, and Marie-Francine Moens. Semeval-2012 task 3: spatial role labeling. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, SemEval '12, pages 365–373, USA, 2012. Association for Computational Linguistics. URL: https://aclanthology.org/S12-1048/.

16 Parisa Kordjamshidi, Marie-Francine Moens, and Martijn van Otterlo. Spatial role labeling: Task definition and annotation scheme. In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10)*, pages 413–420. European Language Resources Association (ELRA), 2010.

17 Klaus Krippendorff. Reliability in content analysis: Some common misconceptions and recommendations. *Human Communication Research*, 30(3):411–433, January 2006. doi:10.1111/j.1468-2958.2004.tb00738.x.

18 J. Richard Landis and Gary G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174, 1977. URL: http://www.jstor.org/stable/2529310.

19 Kevin Lynch. *The Image of the City*. MIT Press, Cambridge, MA, 1960.

20 Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby, and Ben Wellner. SpatialML: Annotation scheme, corpora, and tools. In Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, and Daniel Tapias, editors, *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, May 2008. European Language Resources Association (ELRA). URL: https://aclanthology.org/L08-1017/.

21 Ludovic Moncla, Mauro Gaio, Javier Nogueras-Iso, and Sébastien Mustière. Reconstruction of itineraries from annotated text with an informed spanning tree algorithm. *International Journal of Geographical Information Science*, 30(6):1137–1160, 2016. doi:10.1080/13658816.2015.1108422.

22 Philippe Muller. A qualitative theory of motion based on spatio-temporal primitives. In *Principles of Knowledge Representation and Reasoning*, pages 131–143, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

23 Eric Nichols and Fadi Botros. SpRL-CWW: Spatial relation classification with independent multi-class models. In Preslav Nakov, Torsten Zesch, Daniel Cer, and David Jurgens, editors, *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 895–901, Denver, Colorado, June 2015. Association for Computational Linguistics. doi:10.18653/v1/S15-2150.

24 Michał Olek and Maciej Piasecki. Three-stage extraction of spatial relationships using markers. In Ngoc-Than Nguyen, Bogdan Franczyk, André Ludwig, Manuel Nunez, Jan Treur, Gottfried Vossen, and Adrianna Kozierkiewicz, editors, *Advances in Computational Collective Intelligence*, pages 159–172, Cham, 2024. Springer Nature Switzerland.

25 Ross S. Purves, Stephan Winter, and Werner Kuhn. Places in information science. *Journal of the Association for Information Science and Technology*, 70(11):1173–1182, October 2019. doi:10.1002/asi.24194.

26 James Pustejovsky, Robert Ingria, Roser Saurí, José M. Castaño, Jessica Littman, Robert J. Gaizauskas, Andrea Setzer, Graham Katz, and Inderjeet Mani. The specification language TimeML. In Inderjeet Mani, James Pustejovsky, and Robert J. Gaizauskas, editors, *The Language of Time - A Reader*, pages 545–558. Oxford University Press, 2005.

27 James Pustejovsky, Parisa Kordjamshidi, Marie-Francine Moens, Aaron Levine, Seth Dworman, and Zachary Yocum. SemEval-2015 Task 8: SpaceEval. In Preslav Nakov, Torsten Zesch, Daniel Cer, and David Jurgens, editors, *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 884–894, Denver, Colorado, June 2015. Association for Computational Linguistics. doi:10.18653/v1/S15-2149.

28 James Pustejovsky, Jessica Moszkowicz, and Marc Verhagen. A linguistically grounded annotation language for spatial information. *Traitement Automatique des Langues*, 53(2):87–113, 2012. URL: https://aclanthology.org/2012.tal-2.4/.

29 James Pustejovsky and Jessica L. Moszkowicz. Integrating motion predicate classes with spatial and temporal annotations. In Donia Scott and Hans Uszkoreit, editors, *COLING 2008, 22nd International Conference on Computational Linguistics, Posters Proceedings, 18-22 August 2008, Manchester, UK*, pages 95–98, 2008. URL: https://aclanthology.org/C08-2024/.

30 Hyeong Jin Shin, Jeong Yeon Park, Dae Bum Yuk, and Jae Sung Lee. BERT-based Spatial Information Extraction. In Parisa Kordjamshidi, Archna Bhatia, Malihe Alikhani, Jason Baldridge, Mohit Bansal, and Marie-Francine Moens, editors, *Proceedings of the Third International Workshop on Spatial Language Understanding*, pages 10–17. Association for Computational Linguistics, November 2020. doi:10.18653/v1/2020.splu-1.2.

31 L. Talmy. Lexicalization patterns: Semantic structure in lexical forms. In T. Shopen, editor, *Language Typology and Syntactic Description. Volume III: Grammatical Categories and the Lexicon.* Cambridge University Press, 1985.

32 Thora Tenbrink. Relevance in spatial navigation and communication. In *International Conference on Spatial Cognition*, pages 358–377. Springer, 2012. doi:10.1007/978-3-642-32732-2_23.

33 Barbara Tversky and Paul U. Lee. How space structures language. In Christian Freksa, Christopher Habel, and Karl F. Wender, editors, *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, pages 157–175. Springer, Berlin, Heidelberg, 1998. doi:10.1007/3-540-69342-4_8.

34 Maria Vasardani, Sabine Timpf, Stephan Winter, and Martin Tomko. From descriptions to depictions: A conceptual framework. In Thora Tenbrink, John Stell, Antony Galton, and Zena Wood, editors, *Spatial Information Theory*, pages 299–319, Cham, 2013. Springer International Publishing. doi:10.1007/978-3-319-01790-7_17.

35 Feng Wang, Peifeng Li, and Qiaoming Zhu. A Hybrid Model of Classification and Generation for Spatial Relation Extraction. In Nicoletta Calzolari, Chu-Ren Huang, Hansaem Kim, James Pustejovsky, Leo Wanner, Key-Sun Choi, Pum-Mo Ryu, Hsin-Hsi Chen, Lucia Donatelli, Heng Ji, Sadao Kurohashi, Patrizia Paggio, Nianwen Xue, Seokhwan Kim, Younggyun Hahm, Zhong He, Tony Kyungil Lee, Enrico Santus, Francis Bond, and Seung-Hoon Na, editors, *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1915–1924, Gyeongju, Republic of Korea, October 2022. International Committee on Computational Linguistics. URL: https://aclanthology.org/2022.coling-1.166/.

36 Rudolf Wille. Concept lattices and conceptual knowledge systems. *Computers & Mathematics with Applications*, 23(6):493–515, 1992. doi:10.1016/0898-1221(92)90120-7.

37 Stephan Winter, Ehsan Hamzei, Nico Van de Weghe, and Kristien Ooms. A Graph Representation for Verbal Indoor Route Descriptions. In Sarah Creem-Regehr, Johannes Schöning, and Alexander Klippel, editors, *Spatial Cognition XI*, pages 77–91, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-96385-3_6.

38 Jordan Zlatev. Spatial Semantics. In *The Oxford Handbook of Cognitive Linguistics*. Oxford University Press, 2010. doi:10.1093/oxfordhb/9780199738632.013.0013.

A Appendix: Document Type Definition of refinements and extensions

Document Type Definition (DTD) of refinements and extensions are provided below.

Refinements

<rs:moveLink>

Attributes for <rs:moveLink>

```
<!ELEMENT rs:moveLink EMPTY >
<!ATTLIST rs:moveLink id ID prefix="mvL" #REQUIRED >
<!ATTLIST rs:moveLink relType (traverse | CDATA) #REQUIRED >
<!-- relation type of motion, defaults to traverse or motion class -->
<!ATTLIST rs:moveLink figure IDREF #REQUIRED >
<!-- identifier of the entity that moves along the path -->
<!ATTLIST rs:moveLink ground IDREF #REQUIRED >
<!-- identifier of the relevant place or path -->
<!ATTLIST rs:moveLink source IDREF #REQUIRED >
<!-- identifier of the relevant place or path -->
<!ATTLIST rs:moveLink mids IDREF #REQUIRED >
<!-- identifier of the relevant place or path -->
<!ATTLIST rs:moveLink path IDREF #REQUIRED >
<!-- identifier of the path of movement -->
<!ATTLIST rs:moveLink trigger IDREF #IMPLIED >
<!-- identifier of the motion event that triggered the link -->
<!ATTLIST rs:moveLink sRelation IDREF #REQUIRED >
<!-- identifier of sRelation related to motion event -->
<!ATTLIST rs:moveLink measure IDREF #REQUIRED >
<!-- identifier of the corresponding measure -->
<!ATTLIST rs:moveLink comment CDATA #IMPLIED >
```

▶ **Example 1.** I$_{\text{x1}}$ walk$_{\text{m1}}$ to$_{\text{sr1}}$ the fountain$_{\text{pl1}}$ for [2 minutes]$_{\text{me1}}$.

```
<entity xml:id ="x1" target="I"/>
<motion xml:id="m1" target="walk" motionClass="moveInternal" motionSense="literal"/>
<sRelation xml:id="sr1" target="to" type="goalDefining"/>
<place xml:id="pl1" target="fountain"/>
<measure xml:id="me1" target="2 minutes" value="2" unit="minutes"/>
<rs:moveLink xml:id="mvL1" relType="traverse" figure="#x1" ground="#pl1"
trigger="#m1" sRelation="#sr1" measure="#me1"/>
```

Extensions

<rs:reorientation>

Attributes for <rs:reorientation>

```
<!ELEMENT rs:reorientation (#PCDATA) >
<!ATTLIST rs:reorientation id ID prefix="r" #REQUIRED >
<!ATTLIST rs:reorientation target (IDREF | CDATA) #REQUIRED >
<!ATTLIST rs:reorientation comment CDATA #IMPLIED >
```

▶ **Example 2.** I turn$_{\mathrm{r1}}$ left into Bourke street.

<rs:reorientation xml:id=“r1” target=“turn”/>

<rs:perception>

Attributes for <rs:perception>

<!ELEMENT rs:perception (#PCDATA) >
<!ATTLIST rs:perception id ID prefix=“r” #REQUIRED >
<!ATTLIST rs:perception target (IDREF | CDATA) #REQUIRED >
<!ATTLIST rs:perception comment CDATA #IMPLIED >

▶ **Example 3.** You see$_{\mathrm{pe1}}$ the library.

<rs:perception xml:id=“pe1” target=“see”/>

<rs:reorientLink>

Attributes for <rs:reorientLink>

<!ELEMENT rs:reorientLink EMPTY >
<!ATTLIST rs:reorientLink id ID prefix=“reL” #REQUIRED >
<!ATTLIST rs:reorientLink figure IDREF #REQUIRED >
<!-- identifier of the entity whose orientation changes -->
<!ATTLIST rs:reorientLink ground IDREF #REQUIRED >
<!-- identifier of the referenced place or path -->
<!ATTLIST rs:reorientLink source IDREF #REQUIRED >
<!-- identifier of the referenced place or path -->
<!ATTLIST rs:reorientLink trigger IDREF #IMPLIED >
<!-- identifier of the reorientation event that triggered the link -->
<!ATTLIST rs:reorientLink sRelation IDREF #REQUIRED >
<!-- identifier of sRelation related to reorientation event -->
<!ATTLIST rs:reorientLink comment CDATA #IMPLIED >

▶ **Example 4.** I$_{\mathrm{x1}}$ turn$_{\mathrm{r1}}$ left$_{\mathrm{sr1}}$ into$_{\mathrm{sr2}}$ [Bourke street]$_{\mathrm{p1}}$.

<entity xml:id=“x1” target=“I”/>
<rs:reorientation xml:id=“r1” target=“turn”/>
<sRelation xml:id=“sr1” target=“left”/>
<sRelation xml:id=“sr2” target=“into”/>
<path xml:id=“p1” target=“Bourke Street”/>
<rs:reorientLink xml:id=“reL1” figure=“#x1” ground=“#p1” trigger=“#r1” sRelation=“#sr1, #sr2”/>

<rs:perceiveLink>

Attributes for <rs:perceiveLink>

<!ELEMENT rs:perceiveLink EMPTY >
<!ATTLIST rs:perceiveLink id ID prefix=“peL” #REQUIRED >
<!ATTLIST rs:perceiveLink figure IDREF #REQUIRED >
<!-- identifier of the perceiving entity -->
<!ATTLIST rs:perceiveLink ground IDREF #REQUIRED >
<!-- identifier of the perceived place or path -->
<!ATTLIST rs:perceiveLink trigger IDREF #REQUIRED >
<!-- identifier of the perception event that triggered the link -->
<!ATTLIST rs:perceiveLink comment CDATA #IMPLIED >

▶ **Example 5.** You$_{\mathrm{x1}}$ see$_{\mathrm{pe1}}$ the library$_{\mathrm{pl1}}$.

<entity xml:id="x1" target="You"/>
<rs:perception xml:id="pe1" target="see"/>
<place xml:id="pl1" target="library"/>
<rs:perceiveLink xml:id="peL1" figure="#x1" ground="#pl1" trigger="#pe1"/>