

# Cognitively Sound Route Directions for Web Services

**Editor(s):** Name Surname, University, Country

**Solicited review(s):** Name Surname, University, Country

**Open review(s):** Name Surname, University, Country

Paul Weiser\* and Markus Mayr

*Department of Geodesy and Geoinformation, Research Group Geoinformation,  
Vienna University of Technology, Gusshausstr. 27-29, 1040 Wien, Austria*

**Abstract.** Today's commercial Web routing services provide directions that are mostly quantitative, make no mention of landmarks, and rely on turn-by-turn instructions based on street names. This is in contrast to research findings indicating that people rely on entirely different types of information in the context of route directions. This paper departs from an established theory of cognitively sound invariants in route instructions to propose a formal specification for them. The goal is to provide input for future Web services capable of adapting the type of presented route information to a person's individual needs. We showcase a prototypical application that makes use of our proposed ontology and is capable of presenting route instructions over multiple levels of detail.

**Keywords:** Wayfinding, Route Instructions, Ontology, Cognitive Ergonomics

## 1. Introduction

Route directions on the Web differ substantially from comparable instructions generated by humans. Consider the following two examples of walking directions, guiding you from the Department of Geodesy and Geoinformation in Vienna to a popular café nearby (See Figure 1). Example I was generated by a routing service<sup>1</sup> while example II was generated by a human<sup>2</sup>.

*Example I (Computer-generated)*

1. Head northeast on Gußhausstr. toward Karlsg. (20 m)
2. Slight left onto Karlsg (240 m, 3 min)
3. Slight left toward Treitlstr. (260 m, 3 min)
4. Continue straight onto Treitlstr. Your destination will be on the right (62 m, 54 secs)

*Example II (Human-generated, translated from German)*

1. Go across the street, then turn right at the restaurant
2. Go straight until you have reached a fountain in front of a church.
3. Turn left and pass the university on your left. Continue until you have come to a crosswalk.
4. There, you will see a building that features an owl sculpture - The university library.
5. Go across the street. The destination is about 50 m ahead of you.

Although both examples (most notably instructions 1 through 3, across examples I and II) are semantically "the same" (they lead to the same actions [9]), there are several problems with the instructions in example I.

---

\*Corresponding author. E-mail: weiser@geoinfo.tuwien.ac.at.

<sup>1</sup>Taken from maps.google.com

<sup>2</sup>Thanks to our colleague Jürgen Hahn for providing the directions

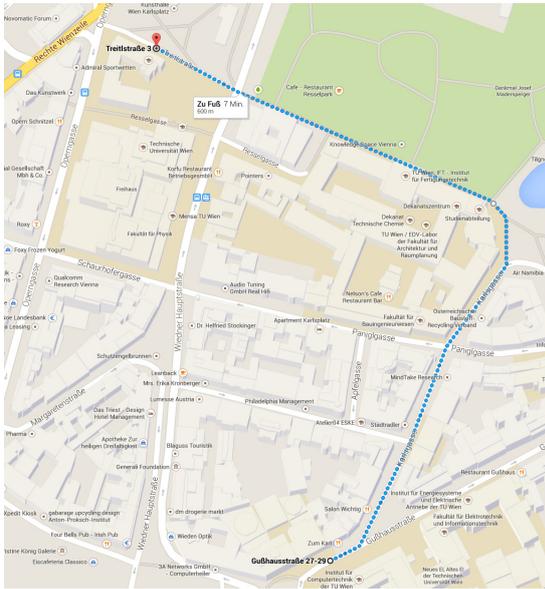


Fig. 1. Example route from the Department of Geodesy and Geoinformation to Café Kunstshalle in Vienna. Source: Google

Computer-generated route information<sup>3</sup> is almost exclusively given in quantitative (metric) terms. Humans, however, rely mostly on qualitative information (cf. [14]). In general, metric information is not very useful unless you have a device to measure it. Most people are not very good at estimating distances to a precision most routing services provide (e.g., turn right after 25 m). People's estimates may not suffice to catch the correct turn and/or resolve situations uniquely with the help of metric instructions (See Figure 2 for two potentially ambiguous situations). Using the GPS device in your automobile works because it gives you a reasonable accurate measurement of distance to the next decision point. This information is periodically updated as you move. Even then, however, it might not be a good idea to constantly check with the navigation system (traffic!) to figure out when exactly one should turn.

Furthermore, instructions of the type "(turn) left toward Treitlstr." can be similarly challenging. First, not every street corner is clearly marked or marked at all. Second, signs might be temporally obstructed, for example, due to construction work. Everyone who has ever tried to desperately find a sign that was not where

<sup>3</sup>We did an evaluation of around 10 commercial routing services and found that none is any different in terms of how information is presented

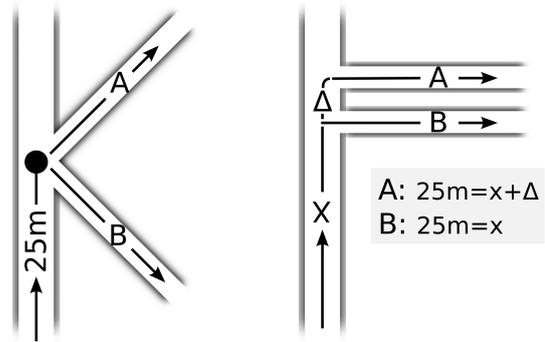


Fig. 2. Two situations that are ambiguous with a quantitative instruction such as "turn right after 25m" ( $\Delta$  denotes the estimation error)

it was supposed to be, knows how unreliable signage can be.

In contrast, the instructions in example I (human-generated) are of entirely different nature. First, most instructions rely on landmarks (salient features in the environment), instead of street names (cf. "restaurant", "church", "university", and "library"). Since landmarks are relatively static and permanent features of the environment there is no need to continuously verify their existence once they have been identified (as opposed to metric information as a function of the current position). One should note, however, that there are instances of landmarks that are less permanent than others. A graffiti can serve as a salient feature used for orientation but is likely to be removed. Second, qualitative information is provided to identify decision points more easily ("in front of", "ahead"). Such directional information is much easier to process and does not require people to estimate distances. Finally, the destination itself is described in much more detail (cf. [24]) to allow for identification ("cross-walk", "owl sculpture", "university library") as opposed to "your destination will be on the right" in example I. Note that only the final instruction mentions metric information ("50m ahead").

Our observations have been confirmed by research on route directions indicating that instructions generated by humans are easier to process (cognitively sound), because they are qualitative [14], include landmarks [17,20,18], and add descriptive information, e.g., hints for self-orientation [7,6,22,3]. See [8,4] for a recent treatment of challenges related to the qualitative and quantitative nature of information generated by humans and computers, respectively, and how they might be solved.

It becomes obvious that there is a discrepancy between what commercial routing services provide today and what can be considered a cognitively sound route instruction. So far there is no "shared understanding" [10] of the structure of route directions that can be used as an input for more adequate Web routing services. The contribution of this paper is three-fold. First, we aim at clarifying some confusing terminology by proposing a taxonomy that unifies several classification approaches of invariants in route instructions. Second, we provide an ontology of way-finding instructions that is grounded in empirical research findings on the nature of cognitively sound instructions. Finally, we demonstrate the benefits of our approach by showcasing a prototypical application that makes use of our ontology and is capable of presenting route instructions over multiple levels of detail.

The remainder of the paper is structured as follows. The next section provides a brief literature review of relevant research findings on route directions. Section 3 reviews existing classifications of invariants in route instructions and proposes a new and unifying taxonomy. Based on that, section 4 describes our cognitively sound ontology of way-finding instructions. Section 5 sketches potential use cases and improvements over common route instructions by showcasing a prototypical application. Section 6 concludes with a summary and future research directions.

## 2. Related Work

### 2.1. The Nature of Communicating Routes

The nature of human language enforces a strict order on how things can be described. If we want to make statements about the world, we have to decide what comes first and what comes next, i.e., we need to map a 3-dimensional (the world) onto a 1-dimensional structure (speech). This issue has become known as "linearization problem"[16]. In case we describe a route, however, we are confined by the linear street-network as part of the discourse, i.e., we can map 1-dimensional route directions directly onto 1-dimensional speech ("first go straight", "then turn right", "then ..."). This fact makes route instructions particularly interesting for research concerning the structure of spatial instructions and how they are communicated by humans.

It is generally acknowledged that direction giving consists of a cognitive and a linguistic phase [6,27]. The former consists of the activation of a mental model

of the geographic space in question, the planning of a route from start to goal, and the selection of features that the instructions should include. The latter consists of the verbal or written output (i.e., the actual utterance) of a route.

The production of route instructions was first studied by linguists (e.g., [27,13]) who analyzed the exchange of spatial instructions in a typical interactive communication setting, e.g., asking someone for directions on the street. Others took an interest in the general structure of route instructions [6,7,1]. This led to the proposal of various invariants humans typically use to communicate route instructions (See Section 3.1 for a more in-depth discussion). Recent research focused on issues related to human-computer interaction, e.g., how computer-generated instructions can be improved and dynamically modified by the user. For example, it has been shown that a person's activity at hand (context) determines the relevancy of a given piece of information. This in turn defines the usefulness of a given route instruction [12] or cartographic map [11] for a particular situation .

It is worthy to note that two different messages (e.g., route instructions), although different in size and/or content may communicate the same message (cf. examples I and II), i.e., they have the same pragmatic semantics [9]. One should keep in mind, however, that messages that are semantically "the same" may not be equally well processed, i.e., they might not be cognitively sound. In general, it is assumed that quantitative information needs to be transformed into its qualitative counterpart (cf. [15,14]). This in turn leads to an increased cognitive workload. In the context of route instructions it has also been shown that humans often group several instructions into one abstract concept, i.e., the spatial information is conceptualized in qualitative terms [14].

### 2.2. What Constitutes Good Route Instructions?

If humans communicate route instructions, the level of detail (LoD) at which information is presented varies and depends, amongst others, on the activity at hand [12]. Computer-generated route instructions have been criticized because they cannot dynamically adjust the presented features of a selected route if the context changes. For example, if someone selects the "most scenic" route they will likely be interested in other features than if somebody selected the fastest route (sight-seeing vs. traffic information). Humans on the other hand are likely to emphasize different features if they

communicate a route to different people and in different contexts.

Also, the "correct" amount of information needed by either participant of a conversation concerning the exchange of route instructions is negotiated or grounded during communication [2]. This happens with the help of signals (e.g. "I do not accept the information you presented") and "mentalizing" the other person's state of knowledge (e.g., "I think I know that you know X") [25]. For example, Tenbrink and Winter [23] highlighted the fact that humans are capable of adjusting the LoD in a "consistent and coherent way, adapting to the addressee's assumed asymmetric information needs". Also, Richter et al. [21] proposed a method to let humans adjust the detail in route directions generated by computers via dialog.

Research has long noted the importance of landmarks as an integral part of human-generated route directions. For example, Denis et al. [7] showed that landmarks are used to solve orientation problems at critical decision points. Also, Lovelace et al. [17] and Michon et al. [18], identified elements, most notably landmarks, that are required for "good route descriptions". Furthermore, it has been suggested that landmarks may be essential in constructing a mental representation of geographic space [18]. To account for their significant role, Raubal and Winter [20] devised a method to automatically extract landmarks based on attribute values from datasets to enrich route instructions.

With this in mind, we can conclude that cognitively sound route instructions should offer different levels of detail, allow for interaction with the presented content and mention landmarks, all of which commercial systems fail to deliver. In this paper, we propose a formal specification of route directions, Web services can use to provide more natural and easily processable route instructions.

### 3. A Taxonomy For Route Instructions

#### 3.1. Existing Classifications

In an attempt to devise a general framework for the quantitative analysis of route instructions Denis [6] analysed a corpus of route instructions produced by humans. He found that participants made consistent use of two main components, landmarks and (prescribing) actions. Based on the main components De-

nis identified five classes into which instructions can be divided.

The first class consists of prescribing actions without reference to a landmark (e.g., "go straight"). The second class contains prescribing actions with reference to a landmark (e.g., "pass the university building"). The third class introduces landmarks without reference to a specific action (e.g. "There is a bank"). The fourth class describes landmarks based on their non-spatial attribute properties (e.g., "The name of the bank is Erste Bank"). The fifth and final class holds comments (e.g., "The route is 15 min. long"). Also, Denis [6] provides an analysis of the content of each class. For example, he mentions that members of the first class either consist of actions that instruct the way-finder to proceed (e.g., "go straight") or to execute a reorientation (e.g., "turn right"). Furthermore, Denis mentions that the third class consists of instances that make explicit references to a spatial location (e.g., "There is a house to your right").

Lovelace et al.[17] in their attempt to identify elements in "good route directions" proposed a finer-grained distinction of landmark features. Their first class consists of landmarks at choice points, i.e., landmarks at which the way-finder turns (e.g., "turn right at church"). The second class consists of on-route landmarks located not at a choice point (e.g., "pass the large green building"). The third and final class includes distant but visible (off-route) landmarks (e.g., "walk towards the tall tower").

In a more recent study, Schwering et al. [22] introduced four additional classes meant to supplement Denis' [6] original classification. The first class includes (self-) orientations using a local landmark (e.g., "behind the church"). The second class consists of (self-) orientations using a global landmark (e.g., "you are looking in the direction of a tower"). The third class includes turning movements using a local landmark (e.g., "turn right at the bank"). The fourth and final class holds non-turning movements using a local landmark (e.g., "follow the street"). Note that streets can be considered as (linear) landmarks.

#### 3.2. Proposed Reclassification

We propose that the different classifications of invariants mentioned in section 3.1 can be merged into a simpler taxonomy. Our unification rests upon the primary distinction between actions and descriptions. The emerging classes provide the basis for our ontology discussed in section 4.

### 3.2.1. Actions

We define an action as a procedure that changes an agent's location on a street network. We take up Denis' [6] distinction between actions that either tell the way-finder to proceed or to reorient. This distinction is reflected in two subclasses for actions, progression and reorientation. A progression keeps an agent's orientation constant while a reorientation changes the direction an agent faces. Additionally, both types of actions may or may not contain references to landmarks. It is possible to make a distinction between local and global landmarks for progressions while a reorientation action always happens at a local landmark (See Section 4). Combining this information (See also Figure 3), we end up with 4 different action type combinations:

1. **ProgressionAtLM**: Progression with reference to a (local/global) landmark: "Pass the university"
2. **ProgressionNoLM**: Progression without reference to a landmark: "Go straight"
3. **ReorientationAtLM**: Reorientation with reference to a (local) landmark: "Turn right at the library"
4. **ReorientationNoLM**: Reorientation without reference to a landmark: "Turn right"

Using this reclassification, Schwering et al.'s [22] class 3 ("turning movement using a local landmark") and Lovelace et al.'s [17] class 1 ("landmark at choice points") can be subsumed in class **ReorientationAtLM**. Furthermore, Schwering et al.'s class 4 ("non-turning movement using local landmark") and both Lovelace's class 1 ("on-route landmarks") and class 3 ("off-route landmarks") can be subsumed in class **ProgressionAtLM**.

### 3.2.2. Descriptions

We define a description as a procedure that updates an agent's knowledge about the world. This can include information about his current position, attributes concerning visible landmarks, or meta-information (e.g., total duration) about the route.

Descriptions have an informative character and can either be of type locating or non-locating. Locating descriptions help the way-finder to self-locate by explicitly referring to a landmark. Locating landmark descriptions can or can not include references to a spatial location. An example for the former is "There is a bank to your left" while an example for the latter is "There is a bank". The locating character is given implicitly, i.e., if the agent can identify the feature mentioned in

the description she can update her knowledge about the current position.

In contrast, non-locating descriptions can either provide additional attributive information on specific landmarks (e.g., "The building is green and called Postbank") or include meta-information on the entire route or route sections (e.g., "The route is 15 minutes long"). It is possible to make a distinction between local and global landmarks for locating and non-locating descriptions. The following combinations (See also Figure 3) are possible:

1. **LocatingAtLMSpatialRef**: Locates the way-finder by making an explicit spatial reference to a (local/global) landmark: "There is a bank to your left"
2. **LocatingAtLMNoSpatialRef**: Locates the way-finder by mentioning a visible (local/global) landmark but without explicitly referring to its spatial location: "There is a bank".
3. **NonLocatingAtLM**: Mentions attributive information of a (local/global) landmark: "The bank has a greenish façade".
4. **Meta**: (Meta-)comment about the route: "The route is 15 minutes long".

With our new classification, Schwering et al.'s [22] class 1 and 2 ("self-orientations") are subsumed in class **LocatingAtLMSpatial**. Denis' [6] class 4 ("landmark attribute information") can be subsumed in class **NonLocatingAtLM** while class 5 is subsumed in **Meta**. Figure 3 shows the possible combinations of actions and descriptions and their subclasses as a function of a reference to landmarks.

## 4. Route Instruction Ontology

This section builds on the proposed reclassification of the invariants found in route directions communicated by humans (See Section 3.2). The ontology was formalized in OWL using Protégé<sup>4</sup>. For the complete definition please refer to the on-line version<sup>5</sup>

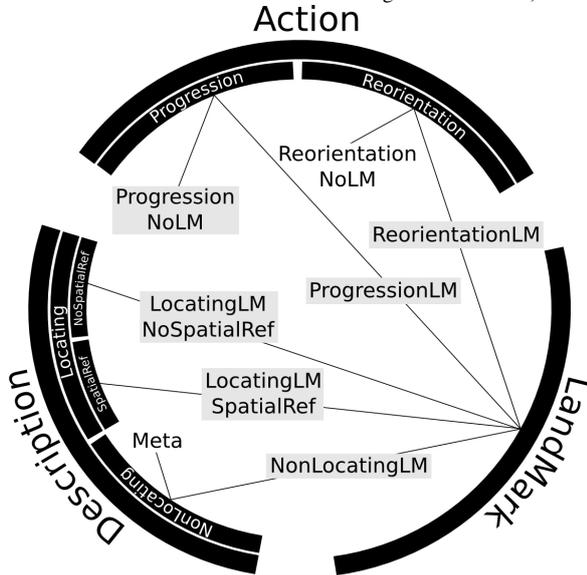
### 4.1. Classes (Excerpt)

We defined three main classes to reflect the primary distinction between actions and descriptions as well as their possible reference to landmarks. The follow-

<sup>4</sup><http://protege.stanford.edu/>

<sup>5</sup><http://goo.gl/DRbT22>

Fig. 3. Proposed Reclassification (for the sake of readability the figure makes no distinction between local and global landmarks)



ing code snippet shows the definition for the class "description" (action and landmarks are defined in a similar fashion):

```
Class: Description
```

```
Description EquivalentTo
  (LocatingDescription or
   NonLocatingDescription)
Description SubClassOf Thing
```

```
ObjectProperty: hasDescription
```

```
hasDescription Range Description
Functional: hasDescription
```

```
Class: LocatingDescription
```

```
LocatingDescription EquivalentTo
  (NoSpatialRef or
   SpatialRef)
LocatingDescription SubClassOf
  Description
LocatingDescription DisjointWith
  NonLocatingDescription
```

The following definition defines an atomic route instruction to either consist of an action or description.

```
Class: RouteInstruction
```

```
RouteInstruction EquivalentTo
  (hasDescription some Description)
  or (hasAction some Action)
```

To link the route instructions to the physical features of a street graph we defined a custom route segment based on nodes and edges (The existence of a graph ontology to be linked with our definition of route segments is assumed). Route segments can either connect two edges via a common node (edge connector) or two nodes via a common edge (node connector). See section 5.2.1 for a more in-depth discussion. In our ontology, route segments are defined as follows:

```
Class: RouteSegment
```

```
RouteSegment EquivalentTo
  NodeConnector or EdgeConnector
  hasNode some Node
  and (hasEdge some Edge)
  and (hasRouteInstruction
       some RouteInstruction)
```

```
Class: EdgeConnector
```

```
NodeConnector DisjointWith
  EdgeConnector
EdgeConnector SubClassOf
  hasNodeCount exactly 1 Literal
  hasEdgeCount exactly 2 Literal
```

```
Class: NodeConnector
```

```
NodeConnector DisjointWith
  EdgeConnector
NodeConnector SubClassOf
  hasNodeCount exactly 2 Literal
  hasEdgeCount exactly 1 Literal
```

To allow for testing we defined various example route instructions that reflect the taxonomy proposed in Section 3.2. For example, the following code shows the `ProgressionLMLocal` class:

```
Class: ProgressionLMLocal
```

```
ProgressionLMLocal SubClassOf
  hasLandmark only LocalLandmark
ProgressionLMLocal SubClassOf
  NamedRouteInstruction
  hasAction only ProgressionAction
```

```

hasAction some ProgressionAction
hasLandmark some LocalLandmark

```

#### 4.2. (Inferred) Model and Testing

We tested the ontology by creating individuals corresponding to an "atomized version" of the human-generated route instruction mentioned in Section 1. For example, the first part of the instruction "Go across the street, then turn right at the restaurant" can be split up in (1) GoAcrossStreet (2) ThereIsARestaurantInFrontOfYou and (3) TurnRightRestaurant. Thus, the individuals are instances of the classes Progression-AtLMLocal, LocatingLMLocalSpatialRef and ReorientationLMLocal, respectively. The internal reasoner provided by Protege was used to test the correct classification of both individuals and classes of a number of possible route instructions. See Figure 4 for the asserted model and Figure 5 for the inferred model. For the complete specification we refer to the ontology file published on-line<sup>6</sup>.

### 5. Prototypical Application

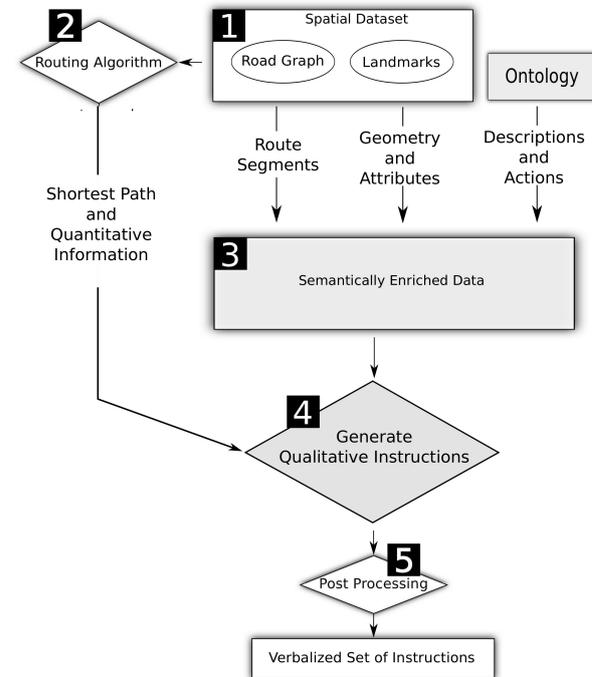
This section sketches a work-flow how cognitively sound route instructions can be generated by integrating existing datasets and applying our proposed ontology. Using examples, we demonstrate how a future Web service could offer routing instructions to a user who wants to get from the Department of Geodesy and Geoinformation to Café Kunsthalle in Vienna (See also Section 1).

#### 5.1. Work-flow

Figure 6 proposes the components a Web service capable of delivering cognitively sound route instructions needs to integrate. First, a *spatial dataset* is required that contains a graph representation of the street network suitable for routing. A routing algorithm (cf. [5]) can utilize this dataset to calculate the shortest path between start and goal of a route. This is the minimum requirement to provide quantitative step-by-step instructions of the type seen in today's Web routing services (cf. example I in Section 1).

To realize the notion of cognitively sound route instruction, landmark information, i.e., their geome-

Fig. 6. Structure of the work-flow (Gray parts are the focus of this work)



try and attributes needs to be integrated. Landmarks can be extracted from already existing datasets by using the method suggested by Raubal and Winter [20]. For example, a suitable dataset offering various information regarding landmarks is provided by OpenStreetMap (OSM) [19].

As a next step, a semantically enriched dataset needs to be created. This requires a method to merge route segments (See 5.2.1) and landmark information into a database that adheres to our proposed ontology of cognitively sound route instructions. The nature of this dataset is described in more detail in section 5.2.

Furthermore, a procedure is required that generates the actual qualitative route instructions. This step needs to combine the quantitative information provided by the routing algorithm and the semantically enriched dataset. The output are qualitative instructions in the form of actions and/or descriptions. The principles of this step are discussed in Section 5.3.

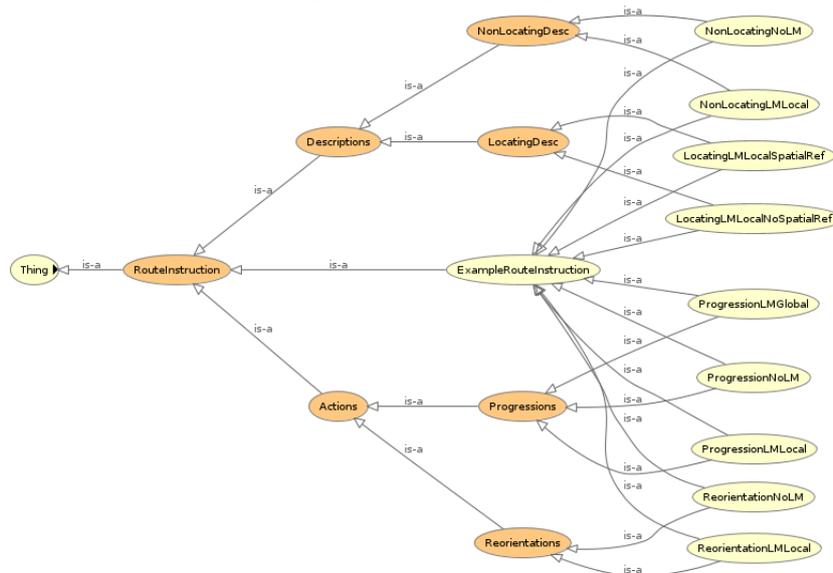
Finally, a post-processing step is needed to let the qualitative instructions appear more like human verbal utterances. In this step the qualitative set of instructions is transformed into a natural language representation. Also, possible redundant parts of the generated routing instructions are removed and multiple single-ton statements combined into one sentence (cf. [14]).

<sup>6</sup><http://goo.gl/DRbT22>

Fig. 4. Asserted Model (Protege)



Fig. 5. Inferred Model (Protege)



## 5.2. Semantically enriched routing data

### 5.2.1. Route-Segments

We propose route segments ("rs") as the data-structure to store qualitative instructions. A route-segment connects either two nodes or two edges in a graph representing a street network. This can be done in two ways:

- Either, by specifying a node indicating the start of a segment, then naming the connected edge and finally adding a node to represent the end of the segment.

- Or, by specifying an edge indicating the start of the segment, crossing a connected node and finally ending at another edge as the final element.

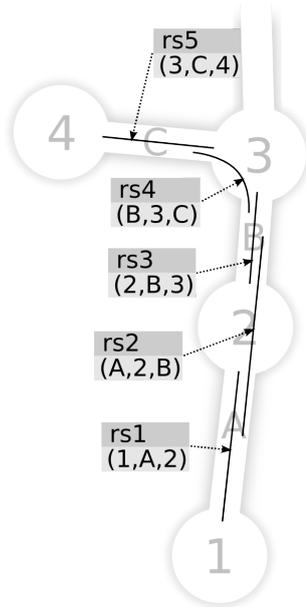
Formally, a route-segment can be described as follows:

$$\text{route-segment} = \begin{array}{l} (\text{Node}, \text{Edge}, \text{Node}) \\ | \\ (\text{Edge}, \text{Node}, \text{Edge}) \end{array}$$

The benefit of using route segments is that one can operate on a single data structure. For example, when describing a landmark in relation to a particular segment its description is valid for the complete segment. It is not necessary to distinguish between nodes and

edges explicitly, as opposed to dealing with common graph representations. For each route segment there is at least one associated valid instruction.

Fig. 7. Five exemplary route-segments (rs1,rs2,rs3,rs4,rs5) for the route from 1 to 4



For example, the route from node 1 to 4 (see Figure 7) can be described as a list of route-segments along the agent travels:

```
(1, A, 2)      == rs1
(A, 2, B)     == rs2
(2, B, 3)     == rs3
(B, 3, C)     == rs4
(3, C, 4)     == rs5
```

Note, that there are two possible combinations for every route-segment. They represent the two corresponding directions an agent can pass a segment. For example, segment "rs4" in Figure 7, defined by "(B,3,C)", is different from "(C,3,B)". Both may share the same space but have to be traveled in opposite directions. This information is needed because spatial relations are expressed in relation to the facing direction of an agent (e.g., "the bank is to your left"). Thus, an intersection with two roads is defined by 8 route-segments, two for each pair of streets.

### 5.2.2. Storage

Using our proposed ontology a database can be populated that stores the semantically enriched data including route segments and landmark information. Note that this paper does not propagate a particular database paradigm. An application as proposed here could be based on both relational and non-relational (nosql) databases. Therefore, the tables for descriptions (see Table 1) and actions (see Table 2) should not be misunderstood as adhering to a particular approach.

The dataset for *descriptions* contains information (e.g., attributes) regarding landmarks relative to a route-segment (See Table 1). Note that the table only displays segments relevant to the example route in gray (See Figure 9)). For each route-segment, entries are computed in advance, producing at least qualitative instruction. The database is populated with all statements that can be computed for a given route segment. For example, the entry "(I,5,M)" in Table 1 contains information regarding a landmark of type "building" that is located on the "left" side of the directed route-segment and has two attributes, namely "has owl sculpture" and "is named library". If there are multiple attributes available for a particular route segment, instructions of different levels of detail (LOD) can be generated from the same database. In fact, the different types (or classes) of the ontology generate a lattice of route instructions [26]. Figure 8 displays possible representations that can be generated from the entry at route segment (I,5,M). For example, "There is a building" (LocatingLMNoSpatialRef) is more general than "There is a building to your left" (LocatingLMSpatialRef) which is more general than "There is a building to your left and it is named library".

The algorithm generating the qualitative instructions (See Section 5.3) decides on the correct LoD of the representation depending on the criteria specified by the user up front. If a user later decides that more or less information is needed for a given route segment the semantically enriched dataset allows for an adjustment of the LoD dynamically (e.g., inclusion/exclusion of spatial references)

The dataset for *actions* (see Table 2) describes how an agent should advance from one route segment to the next. Instructions can include an optional reference to landmarks. For example, at route-segment "(H,4,I)" it is necessary to "turn left at" the "fountain" to get from "H" to "I". As with the dataset for descriptions, every possible route-segment is computed in advance and each entry may contain multiple representations at different LoD for one route-segment. As with descrip-

Table 1

Database for Descriptions, the bold entry is explained in more detail in Figure 8

route-segment	landmark	locating	attribute
( T , 1 , C )	restaurant	in front	
( C , 2 , D )	bike shop	left	[small]
( H , 4 , I )	fountain	right	[big]
( 4 , I , 5 )	university	left	[named TU Vienna]
( I , 5 , M )	crosswalk		
<b>( I , 5 , M )</b>	<b>building</b>	<b>left</b>	<b>[has owl sculpture , is named Library ]</b>
( 5 , M , target )	cafe	right	

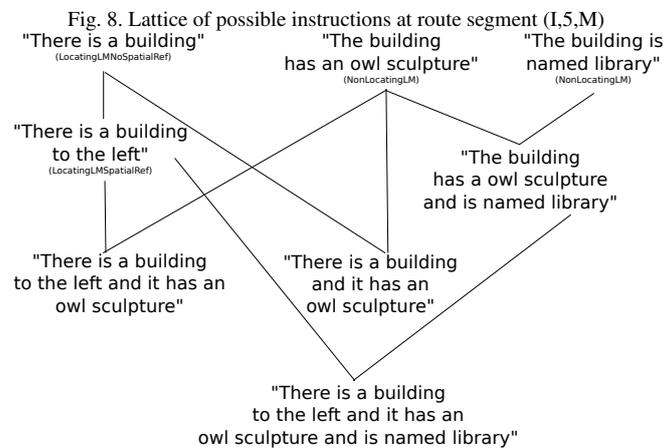


Table 2

Database for actions, bold entry is explained in detail in Table 3

route-segment	landmark	action
( start , T , 1 )	street T	go straight
( T , 1 , C )	restaurant	turn right at
( 1 , C , 2 )		go straight
( C , 2 , D )	bike shop	go straight at
( 2 , D , 3 )		go straight
( D , 3 , H )		go straight
( 3 , H , 4 )		go straight
<b>( H , 4 , I )</b>	<b>fountain</b>	<b>turn left at</b>
( 4 , I , 5 )	university	pass
( I , 5 , M )	building	go straight at
( 5 , M , target )	cafe	arrive at

Table 3

Possible interpretations for the bold formatted entry in Table 2

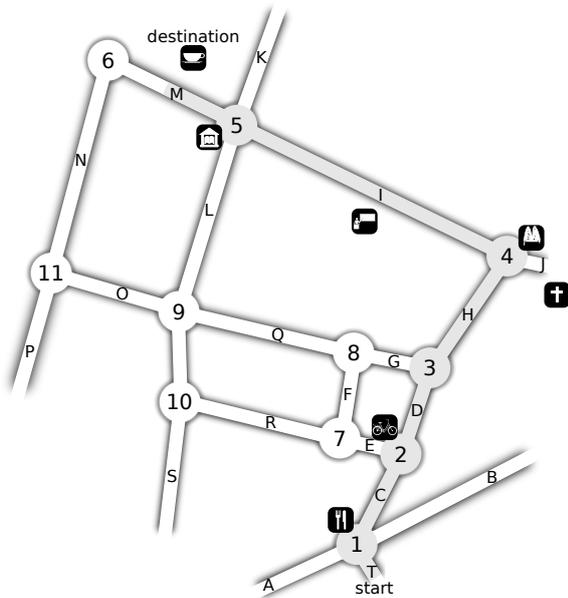
route-segment	landmark	action	type
<b>( H , 4 , I )</b>	<b>fountain</b>	<b>turn left at</b>	
		turn left	ReorientationNoLM
	fountain	turn left at	ReorientationLM(Local)

tions, the algorithm generating a set of qualitative routing instructions may decide on the particular LoD for an action presented to the user. Actions also generate a lattice of possible combinations. How these might be interpreted by a database can be seen in Table 2. Note that the column "type" refers to the corresponding class in the ontology.

### 5.3. Generating Qualitative Routing Instructions

To demonstrate the functionality of our prototypical application, qualitative instructions (see Table 4) for the example route in Figure 9 are calculated. While the method shown here is capable of presenting instructions at different LoD we assume the user requested the most detailed information available for each route segment.

Fig. 9. The example route in gray used for demonstration. Nodes and edges are denoted by numbers and letters, respectively. Landmarks along the route are indicated by their corresponding symbols.



It is worthy to mention that the start and end points of a route require special attention. Depending on where the actual points are, route-segments specifically designed to map these points might have to be calculated on-the-fly. In case of the exemplary route discussed here, this can be seen by the first and last route-segment in Table 2. Segments "(start,T,1)" and "(5,M,target)" were not part of the underlying graph representation of the street network and thus could

not have been calculated prior to the routing request. Both segments, however are needed to make sure the agent reaches the destination even if its corresponding point does not coincide with an already existing route-segment.

Table 4 demonstrates how each route-segment along the route matches a corresponding entry for descriptions and/or actions. For example, at "(H,4,I)", two descriptions are available, namely "(there is a) fountain to the right" and "(the) fountain is big". Since there is also the action "turn left at" ("ReorientationLM(Local)") available for the corresponding landmark "fountain" the instruction for this route segment can be generated by combining both descriptions and the action to yield: "There is a fountain to the right and it is big. Turn left at the fountain". If a user wants less detail, and for example does not need the landmark information to navigate the instruction "turn left" ("ReorientationNoLM") could be presented instead.

At the next route-segment "(4,I,5)", the descriptions "university to the left" and "university is named TU Vienna" as well as the action "pass university" are selected from the tables for descriptions and actions, respectively. This can be combined to "There is a university to the left and it is named TU Vienna. Pass University".

In case of segment "(I,5,M)" which is close to the destination, four different descriptions are available. Because we assume a user who wants to retrieve the highest LoD available the instruction at this particular route segment mentions a crossway ("LocatingLM(Local)NoSpatialRef") and a building at the left ("LocatingLM(Local)SpatialRef"). The building is then further specified by mentioning two "NonLocatingLM(Local)" descriptions, namely telling the user that the building's name is library and that it features an owl sculpture. Also, at route-segment "(I,5,M)" there is multiple landmark information available (the crosswalk and the building named library). Note that the corresponding entry for "action" refers to the building which is named after introducing the landmark "crosswalk". This shows that it is necessary for the system to make sure that any landmarks referenced in the column "action" are first introduced in the field for description.

As a final step, the set of instructions displayed in Table 4 may be post-processed to remove redundancies and let the instructions sound more like human speech. This generates an output similar to Table 5. For example, the instructions for "(C,2,D)", "(2,D,3)", "(D,3,H)" and "(3,H,4)" are summarized into "Go ahead at the

Table 4  
Qualitative Route Instruction (reading instructions: each row from left to right)

	route-segment	description	action
S	( start , T , 1 )		go straight at street T
↓	( T , 1 , C )	restaurant in front	turn right at restaurant
↓	( 1 , C , 2 )		go straight
↓	( C , 2 , D )	bike shop to the left	go straight at bike shop
↓	( 2 , D , 3 )		go straight
↓	( D , 3 , H )		go straight
↓	( 3 , H , 4 )		go straight
↓	( H , 4 , I )	fountain to the right fountain is big	turn left at fountain
↓	( 4 , I , 5 )	university to the left university is named TU Vienna	pass university
↓	( I , 5 , M )	crosswalk building to the left building is named library building is with owl sculpture	go straight at building
T	( 5 , M , target )	café at right	arrive at café

Table 5  
Verbalization of route

- 
- 1 Start.
  - 2 Turn right at the restaurant in front of you.
  - 3 Continue.
  - 4 Go ahead at the bike shop to your left.
  - 5 Turn left at the big fountain to your right.
  - 6 Pass the university named TU Vienna to your left.
  - 7 You will find a crosswalk. At the building named Library which has an owl sculpture to your left, go straight ahead.
  - 8 The Café (destination) is at your right.
- 

bike shop to your left.". Multiple identical instructions (e.g. "go straight") can be removed or chunked into one statement ([14]) since they do not add to the actual information content of the set of instructions.

In case there is no qualitative information available for a particular route-segment or a user requests more detailed information than is available, it is always possible to fall back to quantitative instructions, i.e., turn-by-turn instructions based on metric information and street names. Since the generation of qualitative instructions is based on the route graph, this information is always available.

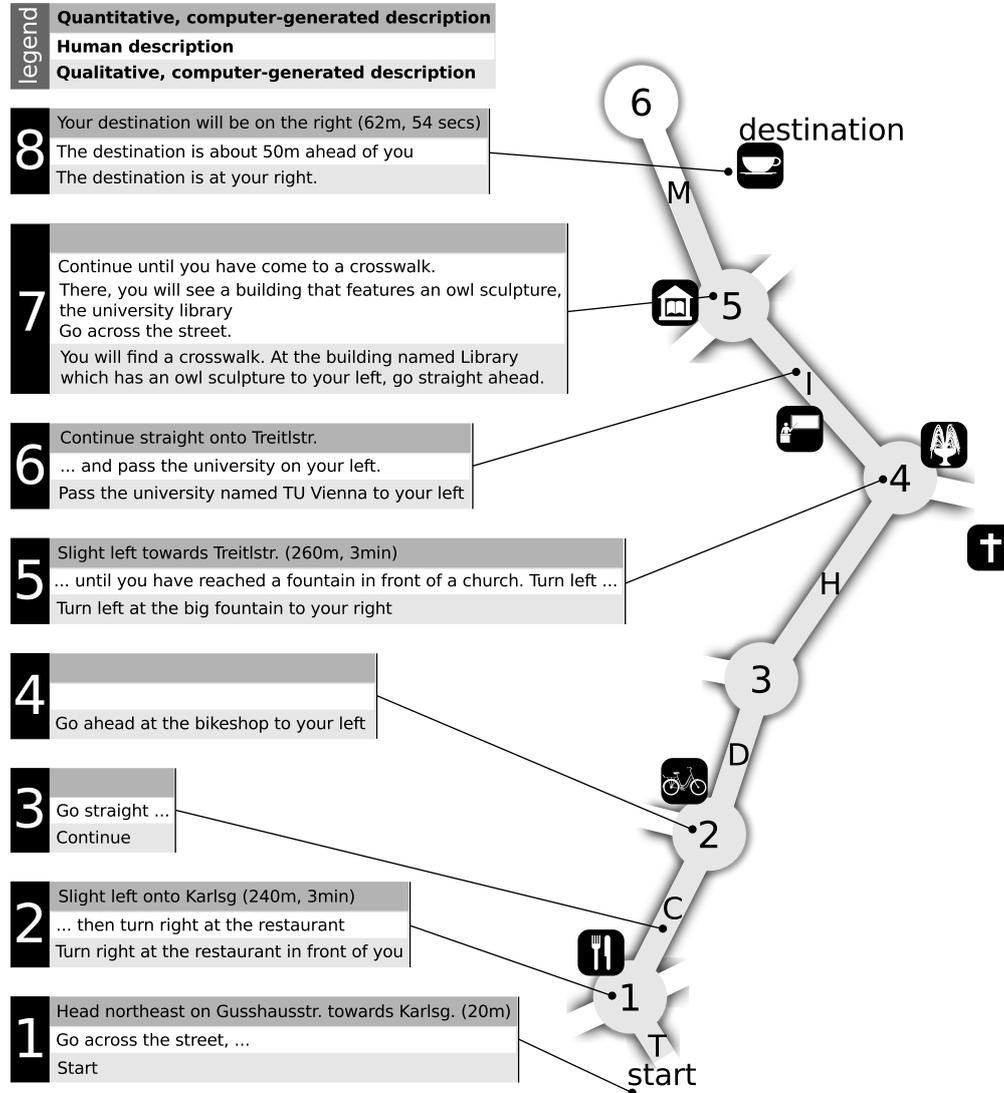
## 6. Conclusion and Future Research Directions

The purpose of this work was to show the discrepancy between how humans produce (and process)

route information and actual implementations of commercial Web routing services. To bridge this gap we proposed to establish a shared understanding of cognitively sound route instructions by developing a simple unified taxonomy and providing a formal specification (ontology). We also demonstrated the benefits of such an approach by showcasing a prototypical application. Our approach allows to generate flexible and dynamic route instructions over multiple levels of detail that resemble human-like instructions. Figure 10 provides a juxtaposition of the three different types of routing instructions discussed in this paper: computer-generated quantitative (typical result of a Web service), human-generated, and computer-generated qualitative (novel approach).

We conclude this paper by mentioning future research directions. First, our ontology can be extended in various ways. For example it would be straight-

Fig. 10. Comparison of three different types of routing information (qualitative computer-generated, human, quantitative computer-generated (landmark-based))



forward to include different forms of spatial reference frames. In this paper we did not make a distinction between intrinsic ("X is to your left") and relative ("X is to the left of Y") reference frames although some people might prefer one form over the other.

Second, semantically aware route instructions can be the basis for more advanced forms of queries. Right now, Web routing services only allow very simple queries of the type "return fastest route from A to B as a function of transport mode T". If a shortest path algorithm operated on a semantic level, our proposed model would yield results to more advanced forms

of queries. For example, one could think of route instructions that include as little reorientation points as possible. Generally, this will not be the fastest route but could minimize the likelihood of getting lost on the way. Also, queries that return additional information regarding landmark attributes (i.e., a special focus on NonLocatingDescription) could be of interest to tourists who want to get historical information regarding landmarks (similar to a city guide). With the widespread availability of Open Street Map (OSM) data [19] in most cities this information is already accessible.

Third, we are planning to define a metric to compare route instructions produced by commercial Web routing services with human-generated instructions. In Section 1 we have already informally discussed some differences. A formal approach, however, can help to make the differences explicit that show between computer-generated and human-generated instructions. The goal is to be able to transform either set of instructions into the respective other representation.

## References

- [1] G. Allen. From knowledge to words to wayfinding: Issues in the production and comprehension of route directions. In S. Hirtle and A. Frank, editors, *Spatial Information Theory A Theoretical Basis for GIS*, volume 1329 of *Lecture Notes in Computer Science*, pages 363–372. Springer Berlin Heidelberg, 1997.
- [2] H. H. Clark. *Using Language*. Cambridge University Press, 1996.
- [3] M.-P. Daniel and M. Denis. Spatial descriptions as navigational aids: a cognitive analysis of route directions. *Kognitionswissenschaft*, 7:45–52, 1998.
- [4] G. De Felice, P. Fogliaroni, and J. O. Wallgrün. A hybrid geometric-qualitative spatial reasoning system and its application in gis. In *Proceedings of the 10th International Conference on Spatial Information Theory*, COSIT'11, pages 188–209, Berlin, Heidelberg, 2011. Springer-Verlag.
- [5] D. Dellling, P. Sanders, D. Schultes, and D. Wagner. Algorithmics of large and complex networks. chapter Engineering Route Planning Algorithms, pages 117–139. Springer-Verlag, Berlin, Heidelberg, 2009.
- [6] M. Denis. The description of routes: A cognitive approach to the production of spatial discourse. *Current Psychology of Cognition*, 16(4):409–458, 1997.
- [7] M. Denis, F. Pazzaglia, C. Cornoldi, and L. Bertolo. Spatial discourse and navigation: an analysis of route directions in the city of venice. *Applied Cognitive Psychology*, 13(2):145–174, 1999.
- [8] P. Fogliaroni. *Qualitative Spatial Configuration Queries - Towards Next Generation Access Methods for GIS*. PhD thesis, University of Bremen, 2013.
- [9] A. U. Frank. Pragmatic information content how to measure the information in a route description. In *Perspectives on Geographic Information Science*, pages 47–68. Taylor & Francis, 2003. semantics.
- [10] T. R. Gruber. A translation approach to portable ontology specifications. *KNOWLEDGE ACQUISITION*, 5:199–220, 1993.
- [11] J. Hahn. Select the appropriate map depending on context in a hilbert space model (scop). In *Proceedings of the 7th International Quantum Interacion Conference*, 2013.
- [12] S. C. Hirtle, S. Timpf, and T. Tenbrink. The effect of activity on relevance and granularity for navigation. In *Proceedings of the 10th international conference on Spatial information theory*, COSIT'11, pages 73–89, Berlin, Heidelberg, 2011. Springer-Verlag.
- [13] W. Klein. Wegauskünfte. *Zeitschrift für Literaturwissenschaft und Linguistik*, 9(33):9–57, 1979.
- [14] A. Klippel. Wayfinding choremes. In W. Kuhn, M. Worboys, and S. Timpf, editors, *COSIT*, number LNCS 2825, pages 301–315, 2003.
- [15] G. Lakoff and R. E. Nunez. *Where Mathematics Comes From: How the embodied mind brings Mathematics into Being*. Basic Books, 2001.
- [16] W. J. Levelt. The speaker's linearization problem. *Philosophical Transactions of the Royal Society of London*, B 295:305–315+, 1981.
- [17] K. L. Lovelace, M. Hegarty, and D. R. Montello. Elements of good route directions in familiar and unfamiliar environments. In *Proceedings of the International Conference on Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science*, COSIT '99, pages 65–82, London, UK, UK, 1999. Springer-Verlag.
- [18] P.-E. Michon and M. Denis. When and why are visual landmarks used in giving directions? In D. Montello, editor, *Spatial Information Theory*, volume 2205 of *Lecture Notes in Computer Science*, pages 292–305. Springer Berlin Heidelberg, 2001.
- [19] F. Ramm and J. Topf. *OpenStreetMap: Die freie Weltkarte nutzen und mitgestalten*. Lehmanns Media GmbH, 2010.
- [20] M. Raubal and S. Winter. Enriching wayfinding instructions with local landmarks. In *GIScience 2002*, pages 243–259. Springer, 2002.
- [21] K.-F. Richter, T. M., and S. Winter. A dialog-driven process of generating route directions. *Computers, Environment and Urban Systems*, 32:233–245, 2008.
- [22] A. Schwering, R. Li, and V. J. A. Anacta. Orientation information in different forms of route instructions. In *Short Paper at the 16th AGILE Conference on Geographic Information Science*, 2013.
- [23] T. Tenbrink and S. Winter. Variable granularity in route directions. *Spatial Cognition and Computation*, 9(1):64 – 93, 2009.
- [24] M. Tomko. *Destination Descriptions in Urban Environments*. PhD thesis, Department of Geomatics, The University of Melbourne, 2007.
- [25] P. Weiser and A. U. Frank. Cognitive transactions - a communication model. In T. Tenbrink, J. Stell, A. Galton, and Z. Wood, editors, *Conference on Spatial Information Theory*, number 8116, pages 129–148, 2013.
- [26] P. Weiser, A. U. Frank, and A. Abdalla. Process composition and process reasoning using partial function application. In *Online Proceedings of the 7th Conference on GIScience, Columbus, Ohio, USA*, 2012.
- [27] D. Wunderlich and R. Reinelt. *Speech, Place, and Action*, chapter How to Get There From Here, pages 183–201. John Wiley & Sons Ltd., 1982.