A NLG-based Application for Walking Directions

Michael Roth and Anette Frank

Department of Computational Linguistics

Heidelberg University

69120 Heidelberg, Germany

{mroth, frank}@cl.uni-heidelberg.de

Abstract

This work describes an online application that uses Natural Language Generation (NLG) methods to generate walking directions in combination with dynamic 2D visualisation. We make use of third party resources, which provide for a given query (geographic) routes and landmarks along the way. We present a statistical model that can be used for generating natural language directions. This model is trained on a corpus of walking directions annotated with POS, grammatical information, frame-semantics and mark-up for temporal structure.

1 Introduction

The purpose of route directions is to inform a person, who is typically not familiar with his current environment, of how to get to a designated goal. Generating such directions poses difficulties on various conceptual levels such as the planning of the route, the selection of landmarks along the way (i.e. easily recognizable buildings or structures) and generating the actual instructions of how to navigate along the route using the selected landmarks as reference points.

As pointed out by Tom & Denis (2003), the use of landmarks in route directions allows for more effective way-finding than directions relying solely on street names and distance measures. An experiment performed in Tom & Denis' work also showed that people tend to use landmarks rather than street names when producing route directions themselves.

The application presented here is an early research prototype that takes a data-driven generation approach, making use of annotated corpora

collected in a way-finding study. In contrast to previously developed NLG systems in this area (e.g. Dale et. al, 2002), one of our key features is the integration of a number of online resources to compute routes and to find salient landmarks. The information acquired from these resources can then be used to generate natural directions that are both easier to memorise and easier to follow than directions given by a classic route planner or navigation system.

The remainder of this paper is structured as follows: In Section 2 we introduce our system and describe the resources and their integration in the architecture. Section 3 describes our corpus-based generation approach, with Section 4 outlining our integration of text generation and visualisation. Finally, Section 5 gives a short conclusion and discusses future work.

2 Combining Resources

The route planner used in our system is provided by the *Google Maps API*¹. Given a route computed in Google Maps, our system queries a number of online resources to determine landmarks that are adjacent to this route. At the time of writing, these resources are: *OpenStreetMaps*² for public transportation, the Wikipedia *WikiProject Geographical coordinates*³ for salient buildings, statues and other objects, *Google AJAX Search API*⁴ for “yellow pages landmarks” such as hotels and restaurants, and *Wikimapia*⁵ for squares and other prominent places.

All of the above mentioned resources can be queried for landmarks either by a single GPS

¹ <http://code.google.com/apis/maps/>

² <http://www.openstreetmap.org>

³ http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Geographical_coordinates

⁴ <http://code.google.com/apis/ajaxsearch>

⁵ <http://www.wikimapia.org>

coordinate (using the `LocalSearch` method in Google AJAX Search and web tools in Wikipedia) or an area of GPS coordinates (using URL based queries in Wikimapia and OpenStreetMaps). The following list describes the data formats returned by the respective services and how they were integrated:

- **Wikimapia and OpenStreetMaps** – Both resources return landmarks in the queried area as an XML file that specifies GPS coordinates and additional information. The XML files are parsed using a JavaScript implementation of a SAX parser. The coordinates and names of landmarks are then used to add objects within the Google Maps API.
- **Wikipedia** – In order to integrate landmarks from Wikipedia, we make use of a community created tool called *search-a-place*⁶, which returns landmarks from Wikipedia in a given radius of a GPS coordinate. The results are returned in an HTML table that is converted to an XML file similar to the output of Wikimapia. Both the query and the conversion are implemented in a *Yahoo! Pipe*⁷ that can be accessed in JavaScript via its URL.
- **Google AJAX Search** – The results returned by the Google AJAX Search API are JavaScript objects that can be directly inserted in the visualisation using the Google Maps API.

3 Using Corpora for Generation

A data-driven generation approach achieves a number of advantages over traditional approaches for our scenario. First of all, corpus data can be used to learn directly how certain events are typically expressed in natural language, thus avoiding the need of manually specifying linguistic realisations. Secondly, variations of discourse structures found in naturally given directions can be learned and reproduced to avoid monotonous descriptions in the generation part. Last but not least, a corpus with good coverage can help us determine the correct selection restrictions on verbs and nouns occurring in directions. The price to pay for these advantages is

⁶ <http://toolserver.org/~kolossos/wp-world/umkreis.php>

⁷ <http://pipes.yahoo.com/pipes/pipe.info?id=BB10x8G73RGbWzKnBR50VA>

the cost of annotation; however we believe that this is a reasonable trade-off, in view of the fact that a small annotated corpus and reasonable generalizations in data modelling will likely yield enough information for the intended navigation applications.

3.1 Data Collection

We currently use the data set from (Marciniak & Strube, 2005) to learn linguistic expressions for our generation approach. The data is annotated on the following levels:

- Token and POS level
- Grammatical level (including annotations of main verbs, arguments and connectives)
- Frame-semantics level (including semantic roles and frame annotations in the sense of (Fillmore, 1977))
- Temporal level (including temporal relations between discourse units)

3.2 Our Generation Approach

At the time of writing, our system only makes use of the first three annotation levels. The lexical selection is inspired by the work of Ratnaparkhi (2000) with the overall process designed as follows: given a certain situation on a route, our generation component receives the respective frame name and a list of semantic role filling landmarks as input (cf. Section 4). The generation component then determines a list of potential lexical items to express this frame using the relative frequencies of verbs annotated as evoking the particular frame with the respective set of semantic roles (examples in Table 1).

SELF MOTION	
PATH	17% <i>walk</i> , 13% <i>follow</i> , 10% <i>cross</i> , 7% <i>continue</i> , 6% <i>take</i> , ...
GOAL	18% <i>get</i> , 18% <i>enter</i> , 9% <i>continue</i> , 7% <i>head</i> , 5% <i>reach</i> , ...
SOURCE	14% <i>leave</i> , 14% <i>start</i> , ...
DIRECTION	25% <i>continue</i> , 13% <i>make</i> , 13% <i>walk</i> , 6% <i>go</i> , 3% <i>take</i> , ...
DISTANCE	15% <i>continue</i> , 8% <i>go</i> , ...
PATH + GOAL	29% <i>continue</i> , 14% <i>take</i> , ...
DISTANCE + GOAL	100% <i>walk</i>
DIRECTION + PATH	23% <i>continue</i> , 23% <i>walk</i> , 8% <i>take</i> , 6% <i>turn</i> , 6% <i>face</i> , ...

Table 1: Probabilities of lexical items for the frame SELF_MOTION and different frame elements

For frame-evoking elements and each associated semantic role-filler in the situation, the grammatical knowledge learned from the annotation level determines how these parts can be put together in order to generate a full sentence (cf. Table 2).

SELF MOTION	
<i>walk</i> + [<i>building</i> _{PATH}]	<i>walk</i> → <i>walk</i> + PP PP → along + NP NP → the + <i>building</i>
<i>get</i> + [<i>building</i> _{GOAL}]	<i>get</i> → <i>get</i> + to + NP NP → the + <i>building</i>
<i>take</i> + [<i>left</i> _{DIRECTION}]	<i>take</i> → <i>take</i> + NP NP → a + <i>left</i>

Table 2: Examples of phrase structures for the frame SELF_MOTION and different semantic role fillers

4 Combining Text and Visualisation

As mentioned in the previous section, our model is able to compute single instructions at crucial points of a route. At the time of writing the actual integration of this component consists of a set of hardcoded rules that map route segments to frames, and landmarks within the segment to role fillers of the considered frame. The rules are specified as follows:

- A turning point given by the Google Maps API is mapped to the SELF_MOTION frame with the actual direction as the semantic role *direction*. If there is a landmark adjacent to the turning point, it is added to the frame as the role filler of the role *source*.
- If a landmark is adjacent or within the starting point of the route, it will be mapped to the SELF_MOTION frame with the landmark filling the semantic role *source*.
- If a landmark is adjacent or within the goal of a route, it will be mapped to the SELF_MOTION frame with the landmark filling the semantic role *goal*.
- If a landmark is adjacent to a route or a route segment is within a landmark, the respective segment will be mapped to the SELF_MOTION frame with the landmark filling the semantic role *path*.

5 Conclusions and Outlook

We have presented the technical details of an early research prototype that uses NLG methods

to generate walking directions for routes computed by an online route planner. We outlined the advantages of a data-driven generation approach over traditional rule-based approaches and implemented a first-version application, which can be used as an initial prototype extensible for further research and development.

Our next goal in developing this system is to enhance the generation component with an integrated model based on machine learning techniques that will also account for discourse level phenomena typically found in natural language directions. We further intend to replace the current hard-coded set of mapping rules with an automatically induced mapping that aligns physical routes and landmarks with the semantic representations. The application is planned to be used in web experiments to acquire further data for alignment and to study specific effects in the generation of walking instructions in a multimodal setting.

The prototype system described above will be made publicly available at the time of publication.

Acknowledgements

This work is supported by the DFG-financed innovation fund FRONTIER as part of the Excellence Initiative at Heidelberg University (ZUK 49/1).

References

- Dale, R., Geldof, S., & Prost, J.-P. (2002). Generating more natural route descriptions. *Proceedings of the 2002 Australasian Natural Language Processing Workshop*. Canberra, Australia.
- Fillmore, C. (1977). The need for a frame semantics in linguistics. *Methods in Linguistics*, 12, 2-29.
- Marciniak, T., & Strube, M. (2005). Using an annotated corpus as a knowledge source for language generation. *Proceedings of the Workshop on Using Corpora for Natural Language Generation*, (pp. 19-24). Birmingham, UK.
- Ratnaparkhi, A. (2000). Trainable Methods for Surface Natural Language Generation. *Proceedings of the 6th Applied Natural Language Processing Conference*. Seattle, WA, USA.
- Tom, A., & Denis, M. (2003). Referring to landmark or street information in route directions: What difference does it make? In W. Kuhn, M. Worboys, & S. Timpf (Eds.), *Spatial Information Theory* (pp. 384-397). Berlin: Springer.

Figure 1: Visualised route from *Rohrbacher Straße 6* to *Hauptstrasse 22, Heidelberg*. Left: GoogleMaps directions; Right: GoogleMaps visualisation enriched with landmarks and directions generated by our system (The directions were manually inserted here as they are actually presented step-by-step following the route)

Script Outline

Our demonstration is outlined as follows: At first we will have a look at the textual outputs of standard route planners and discuss at which points the respective instructions could be improved in order to be better understandable or easier to follow. We will then give an overview of different types of landmarks and argue how their integration into route directions is a valuable step towards better and more natural instructions.

Following the motivation of our work, we will present different online resources that provide landmarks of various sorts. We will look at the information provided by these resources, examine the respective input and output formats, and state how the formats are integrated into a common data representation in order to access the information within the presented application.

Next, we will give a brief overview of the corpus in use and point out which kinds of annotations were available to train the statistical generation component. We will discuss which other annotation levels would be useful in this scenario and which disadvantages we see in the current corpus. Subsequently we outline our plans to acquire further data by collecting directions for routes computed via Google Maps, which would allow an easier alignment between the instructions and routes.

Finally, we conclude the demonstration with a presentation of our system in action. During the presentation, the audience will be given the possibility to ask questions and propose routes for which we show our system's computation and output (cf. Figure 1).

System Requirements

The system is currently developed as a web-based application that can be viewed with any JavaScript supporting browser. A mid-end CPU is required to view the dynamic route presentation given by the application. Depending on the presentation mode, we can bring our own laptop so that the only requirements to the local organisers would be a stable internet connection (access to the resources mentioned in the system description is required) and presentation hardware (projector or sufficiently large display).