Data driven learning of the meaning of route descriptions
Project specification

Erik Odenman

April 9, 2013

Coordinator: Gabriel Skantze

1 Introduction

The map task involves creating a program that, given a map and a route description given in natural language, can understand and draw this route. The complexity of the language and the many ways a particular route can be described makes this a very challenging problem, as with most problems involving communication between human and computer via speech. Furthermore many applications of the problem are imaginable, for example in computer games and in the field of robotics just to name a few.

Although this task can be approached using semantics and a long list of spatial words and known landmarks, a language independent and non rule-based approach would be more convenient. Therefore, in this project very little knowledge will be given to the program and a big part of the challenge will be to extract information about the map's landmarks from sample descriptions.

2 Problem statement

Given a map and a description of some path, the task is to follow this path as closely as possible and draw it on the map. The descriptions should be as natural as possible, and although english will be used for this project the algorithms developed shouldn't be limited to this language. To provide data for the program a set of sample descriptions are provided. These descriptions are divided into instructions, where each instruction consists of some raw text as well as coordinates on the map describing where the instruction was uttered.

3 Approach

To make the problem easier to approach I will focus only on the search for landmarks. The side on which landmarks are passed will be ignored as well as the meaning of cardinal directions. This creates a very simple approach that still could work due to the excessive nature of the language.

The first problem that arises is how to recognize the words that describe landmarks, in contrast to unimportant words like 'to'. A technique that is commonly used to decide a word's relevance to a document is *tf-idf* (Term frequency - Inverse Document Frequency). It works by selecting words that occur as much as possible within a given document (term frequency), but at the same time doesn't occur in too many other documents (inverse document frequency). A variant of this technique could be used for the purpose of this project (because words describing landmarks should be common in one part of the map but not elsewhere).

When the most significant word for every instruction is isolated an algorithm to construct a path from these words is needed. For this purpose two different algorithms will be studied, one greedy and one using dynamic programming. These algorithms will then be tested on a subset of the sample descriptions (not the ones used for collecting data) and it will be observed how well the algorithms perform and whether the results differ or not.

4 References

Learning to Follow Navigational Directions, Adam Vogel and Dan Jurafsky
http://nlp.stanford.edu/pubs/spatial-acl2010.pdf

Interpretation of Spatial Language in a Map Navigation Task, Michael Levit and Deb Roy
http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=04200804

A Testbed for Examining the Timing of Feedback using a Map Task, Gabriel Skantze
http://www.speech.kth.se/prod/publications/files/3761.pdf

Tf-idf (Maybe use a different link, although this seems pretty reliable)
http://en.wikipedia.org/wiki/Tf-idf

5 Time plan

3 february - Project specification
4 february - Start literature study
5 mars - Half way meeting
8 april - Implementation of algorithms complete
10 april - Testing of algorithms
12 april - Essay hand in