Project Periodic Report

Grant Agreement number: 270019
Proposal acronym: SPACEBOOK
Project title: Spatial & Personal Adaptive Communication Environment: Behaviors & Objects & Operations & Knowledge
Funding Scheme: EC FP7, Call 6

Date of latest version of Annex I against which the assessment will be made:

Periodic report: 2nd
Period covered: 1-March-12 - 28-Feb-13

Name, title and organization of the scientific representative of the project's coordinator:

Michael Minock
Associate Professor
Umeå University, Sweden

Tel: +46 90 786 6398
Fax: +46 90 786 6126
E-mail: mjm@cs.umu.se
Project website address: `http://spacebook-project.eu/`

1 Publishable summary

1.1 Project context and objectives

SPACEBOOK aims to provide comfortable hands-free, eyes-free support for urban pedestrians exploring and acquiring knowledge of a city, by taking advantage of, and in some cases, extending the state-of-the-art in speech understanding, natural language processing, question-answering, natural language generation and text-to-speech technology. SPACEBOOK is developed as an open source project where progress is bench marked through controlled task-based experiments with real pedestrians in central Edinburgh. The SPACEBOOK project will continue to generate concrete technical and scientific advances for eyes-free, hands-free navigation and exploration systems which will support applications in tourism. It also has potential application in rural contexts, and among user communities such as the visually impaired, the cognitively impaired and the emergency services (tasks that must be unencumbered by invasive technology). While there has been recent industrial interest in voice-based, mobile city information systems (e.g. Apple's Siri system), SPACEBOOK is unique in pioneering the modeling of the pedestrian's field of view to provide situated speech-based dialogue to support both navigation and exploration. No industrial effort to our knowledge is yet attempting this.

In addition to advances in navigation and exploration systems, SPACEBOOK provides a task environment in which more fundamental scientific and technical knowledge will be generated. Specifically we seek to advance the state of the art in:

I. **Model-based approaches to plan generation and recognition.** There are essentially three approaches to building agents: 1.) they can be hand programmed; 2.) they can be learned from data; 3.) they can be *model-based*, arising from a first principles model of the agent's environment, goals, actions, sensors, etc. The model-based approach is characterized by very flexible behavior, but significant scalability challenges. The SPACEBOOK task environment is a real-world example that will help evaluate the merits of the model-based approach and drive algorithmic innovation in generation and recognition of behavior in the model-based approach. Model-based approach for plan generation is being used for dialogue management, while for plan recognition is being used for goal recognition. Both of these threads are implemented in two software modules developed by UPF.

II. **Statistical learning techniques for interaction management.** The determination of when and in what sequence to execute actions in an uncertain, partially-observable environment is a long-standing topic in Artificial Intelligence. One branch of work focuses on using reinforcement learning over *partially observable Markov decision processes* (POMDPs) to decide on a policy that determines the correct action to execute based on the agent's belief state about the world. There are three key issues that complicate this problem: 1.) determining which features of the world should map into a limited set of state variables for the POMDP; 2.) building realistic simulations to feed reinforcement learning with sufficient data; 3.) algorithmic issues that enable POMDPs to scale to larger numbers of state variables and actions. SPACEBOOK provides a challenging task environment in which to test reinforcement learning for POMDPs, where actions are the next communication action of the SPACEBOOK system.

III. **Machine learning of natural language understanding components.** A very active research area lately has been the learning of natural language understanding components that map natural language (e.g. English) to logical expressions in a *meaning representation language* (MRL)

that can drive applications (e.g. query a database, control a DVD player, etc.). Typically the learning problem takes as input a corpus of natural language expressions paired with their corresponding MRL expressions and the output of learning is a component that can map natural language to MRL expressions. We shall drive innovation in this area by defining an MRL adequate to SPACEBOOK's complex, real world task environment, as well as create a semantic parsing corpus for the SPACEBOOK domain. This is a clearly sought after resource for the machine learning community, which is currently using only a small number of semantic parsing corpora based on simple scenarios such as booking airline flights. Moreover our approach to structural learning should extend the state of the art in learning of natural language understanding components.

We have made considerable progress in the creation of the semantic parsing corpus, which will be used as the basis for developing the semantic parsing component itself. The corpus contains over 1,000 utterances collected using "Wizard of Oz" experiments. These are experiments where a human pretends to be the computer, in order to get realistic interactions between a user and the system. The creation of the corpus requires annotation of these utterances using SPACEBOOK's MRL.

The annotation process is a challenging and time-consuming process, requiring not only annotation of the corpus itself but also the creation of a detailed set of guidelines so that an annotator not familiar with SPACEBOOK can still perform the annotation. This process requires a number of iterations in which the guidelines are refined based on the utterances observed in the data and feedback from annotators.

Two people carried out the bulk of the annotation; one of them a researcher on the project and the other a professional freelance linguist. The annotation was carried out over the period of a few months. We will soon make the corpus publicly available with the expectation that semantic parsing researchers around the world will begin to use it.

While the above areas are our declared basic research areas, there are a host of other areas in which project members are generating basic research results, including: (IV) textual descriptions of vista space, (VI) Question Answering (QA) systems over textual resources, (VII) *minimum recursion semantics (Copestake, et al.)* based meaning representation languages, and more.

1.2 Work performed since the project start and main results to date

In the SPACEBOOK project we have developed a speech-driven, hands-free, eyes-free prototype for pedestrian navigation and exploration which we have evaluated through controlled task-based experiments with real pedestrians in central Edinburgh. This achievement was the result of intense design and requirements collection, followed by implementation of individual components, followed by integrating and testing the whole working system. On August 3, 2012, the SPACEBOOK project reached a critical milestone: *we successfully completed an experiment ready prototype for the streets of Edinburgh.* In the ensuing months we have conducted a systematic evaluation of this prototype with real pedestrians and, based on these experiences, have planned for the refinement and enhancement of this prototype in the final phase of the project.

In addition to this prototype and its evaluation in Edinburgh (see YOUTUBE videos linked at our website), we have integrated and tested SPACEBOOK components into systems that we have tested and experimented with in other cities including Stockholm, Umeå and Barcelona. We have made an initial public data release of our WoZ studies – see `http://www.macs.hw.ac.uk/ilabarchive/`

spacebook/login.php (also linked at the project website). This includes WoZ data collected in both Edinburgh and a separate WoZ study conducted in Stockholm. We shall soon publish a semantic parsing corpus for the SPACEBOOK domain based on the WoZ data from our Edinburgh experiments.

Beyond these achievements, much component level work is ongoing that is advancing the state of the art in many critical areas. For example we have investigated the derivation of saliency measures of geographic features in our city model via both web-based text mining and directly from human route directions. We have also pushed QA technology to the limit to support sequences of speech-based queries over document describing entities in Edinburgh. We have also applied machine learning techniques to information extraction over documents to populate structured relations that may be directly queried. Finally we have explored a variety of alternative representations and processing techniques to support visibility determination, route calculation, dialogue management and the generation of natural language descriptions. A key activity in the project is to compare and contrast alternative techniques to solve problems within the SPACEBOOK task domain.

Dissemination Activities

In the first two years of the project, members have published 25 conference and journal papers, with many additional works under review. In addition we have produced and published a series of YOUTUBE videos that explain our project concept, showcase our city model and pedestrian tracker and demonstrate our working systems. These demonstration videos are linked on our project website. The project has already released WoZ data for public download and will release more data as the project completes. Finally, we look forward to an open source software release of a significant number of our software components and integrated systems.

1.3 Expected final results and potential impact and use

In the final year of the project our systems and components will be stabilized and then modified to maximize objective navigation and exploration evaluation metrics. This will lead to concrete findings on how SPACEBOOK like systems can enhance the way in which people navigate and explore the city. In addition we shall continue to pursue and publish findings in our basic research threads. Our commitment to both open source and open access will make our results directly accessible to a wide audience in the scientific, business and open source communities. In turn, such entities will transfer the ideas, algorithms, components and systems in the SPACEBOOK project into tools to help real people better navigate and explore their cities.

1.4 Project website

SPACEBOOK's public website address is http://spacebook-project.eu.