



spacebook-project.eu

Summary of SpaceBook project results

Michael Minock, Johan Boye, Stephen Clark, Morgan Fredriksson,
Hector Geffner, Oliver Lemon, William Mackaness, Bonnie Webber

Distribution: Public

SpaceBook

Spatial & Personal Adaptive Communication Environment: Behaviors & Objects & Operations
& Knowledge
270019 Deliverable 7.2.2

July 10, 2014



Project funded by the European Community
under the Seventh Framework Programme for
Research and Technological Development



CogSys
Cognitive Systems



The deliverable identification sheet is to be found on the reverse of this page.

Project ref. no.	270019
Project acronym	SpaceBook
Project full title	Spatial & Personal Adaptive Communication Environment: Behaviors & Objects & Operations & Knowledge
Instrument	STREP
Thematic Priority	Cognitive Systems, Interaction, and Robotics
Start date / duration	01 March 2011 / 36 Months
Security	Public
Contractual date of delivery	M36 = March 1, 2014
Actual date of delivery	July 10, 2014
Deliverable number	7.2.2
Deliverable title	Summary of SpaceBook project results
Type	Report
Status & version	Final 2.0
Number of pages	10 (excluding front matter)
Contributing WP	7
WP/Task responsible	UMU
Other contributors	Phil Bartie, Tiphaine Dalmas, Bruce Gittings, Jana Götze, Joakim Gustafson, Robin Hill, Srinivasan Janarthanam, An- ders Jonsson, Simon Keizer, Jürgen Königsmann, Xingkun Liu, Johan Mollevik, Andreas Vlachos, Mattias Åsander
Author(s)	Michael Minock, Johan Boye, Stephen Clark, Morgan Fredriksson, Hector Geffner, Oliver Lemon, William Mack- aness, Bonnie Webber
EC Project Officer	Franco Mastroddi
Keywords	

The partners in SpaceBook are:

Umeå University	UMU
University of Edinburgh HCRC	UE
Heriot-Watt University	HWU
Kungliga Tekniska Högskola	KTH
Liquid Media AB	LM
University of Cambridge	UCAM
Universitat Pompeu Fabra	UPF

For copies of reports, updates on project activities and other SPACEBOOK-related information, contact:

The SPACEBOOK Project Co-ordinator:

Dr. Michael Minock

Department of Computer Science

Umeå University

Sweden 90187

mjm@cs.umu.se

Phone +46 70 597 2585 - Fax +46 90 786 6126

Copies of reports and other material can also be accessed via the project's administration homepage,
<http://www.spacebook-project.eu>

©2014, The Individual Authors.

No part of this document may be reproduced or transmitted in any form, or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission from the copyright owner.

Contents

1	Executive Summary	1
2	Concept and Objectives	2
3	Main Scientific and Technological Results	4
4	The potential impact	5
5	Project website	5

1 Executive Summary

The SPACEBOOK project has pioneered the use of speech-based dialogue to support pedestrian navigation and exploration of urban environments. While there have been many new approaches to this problem, ranging from Google glasses, to more routine multi-modal touch screens, to even haptic interfaces, to our knowledge SPACEBOOK is the first system to rely solely on natural language speech and text-to-speech responses to support real-time, in situ navigation and exploration of the urban environment. 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.

The primary hypothesis of the SPACEBOOK project is that such a speech-only interface is feasible and can be used reliably by a wide range of subjects. A secondary hypothesis is that such a set-up could also be rated as more enjoyable and informative than the alternatives. To test both of these hypotheses the SPACEBOOK team built a series of prototype systems which culminated in a single system that was tested with subjects on the streets of Edinburgh. The results of this evaluation (and others), lend support to our hypotheses that such systems can be made reliable and enjoyable. This suggests that the technology laying behind SPACEBOOK could and should be further developed and refined toward large scale development, deployment and use.

In addition to the overarching goal of confirming the value and promise of the SPACEBOOK vision, we have conducted research in a variety of sub-areas where the SPACEBOOK domain serves as a real-world task environment. The objectives, work performed and results in these focused sub-areas are touched on below.

2 Concept and Objectives

Imagine you are visiting an unknown European city. You arrive at the central train station toward night fall. Now you must find your hotel. You feel uncomfortable asking for directions from random people on the street. You are not so skilled at using maps, and in any case you don't want to be bothered looking down or having to read street signs as you navigate. You are lucky to have an Android phone with the SPACEBOOK app that speaks directions into your ear as you *navigate* your way to your hotel. With your eyes free and while pulling your suitcase, you listen as it guides the way – “OK continue another 100 meters, and turn right into the alleyway immediately after passing a Tesco on your right.”

Now after a good night's sleep, you want to explore the city, learning about its history, attractions and other features. As you walk down the street from the hotel, a stunning neoclassical building on your left comes into view. Without pushing any buttons, you simply say “SpaceBook! What is the building on my left?”. SPACEBOOK answers, “That is the National Art Museum. It houses ... (you listen for 20 seconds to the basic introduction)”. Then you interrupt, “when does it close today?”. SPACEBOOK answers, “It closes at 4 o'clock today.” You continue walking. You have plenty of time (and now enhanced capability) to *explore* the city on foot.

Put simply, SPACEBOOK's core objective is 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. In addition SPACEBOOK provides a task environment in which more focused scientific and technical knowledge was generated. Specifically we advanced 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 helps 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 in the project.
- 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 longstanding 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 we tested reinforcement learning for POMDPs, where actions are the next communication action of the SPACEBOOK system.

III. Machine learning of natural language understanding components. An 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 by creating 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 extends the state of the art in learning of natural language understanding components.

IV. Generating textual descriptions and saliency models of objects in field of views In GIS systems, traditional representations of space are Cartesian in the form they represent spatial dimensions of objects. By combining a visibility engine with a rich description of geographic entities, SpaceBook affords a way of describing 'place', focusing on entities that fall within the field of view, both near and far. By incorporating textual descriptions and vernacular regions in the city model, we are able to describe the city in a more intuitive manner that maps more readily to the user's perception of the city. This significantly extends the 'spatial vocabulary' that we might use to describe the field of view.

The combined visibility engine and detailed city model provides a tailored environment in which to explore the efficacy of textual descriptions of 'constellations' of complex geographic objects (text based scene description), in real time (exploring parallel processing and enhancements to spatial statistical algorithms). Such research is fundamental to speech based interaction in highly dynamic environments.

Theoretical topological models applied to vista space are used to describe the absolute and relative positions of places of interest. This involves modeling egocentric views in rapidly changing vista space, and the generation of egocentric object references (referring expressions). These descriptions are evaluated against a number of metrics, in the context of human subject based navigational tasks in order to assess the 'correctness' of these descriptions against a human's actual perception of that urban space.

V. Question Answering (QA) systems over textual resources

In the past decades, evaluations in the QA research community (e.g. TREC QA, CLEF QA) and in the industry (IBM Watson Jeopardy challenge) have shown that automated systems are increasingly more reliable when tasked with answering certain question types, such as factoids and definition questions. More recent research has focused on handling more complex questions and request environments. In Spacebook, QA is provided as a discovery functionality along with navigation. Our focus is on assessing the integration of existing QA techniques in a larger, multi-task and speech-only mobile system, in particular by making use of Spacebook's simulation of the user's field of view. While mobile applications providing search and city exploration exist (e.g. Field Trip, Google Now currently extending to Google glasses), there is little data on the impact of real street environment and multi-tasking on QA behaviour, and how subjects make use of the technology.

3 Main Scientific and Technological Results

The SPACEBOOK project initiated with intense requirements collection, partly guided through Wizard-of-Oz studies – where a human operator played the role of the system. This impacted the form and content of the data and knowledge representations underlying SPACEBOOK. This also impacted the initial design and implementation of SPACEBOOK components. A range of experimental work contributed to the final design of the SPACEBOOK system, the implementation of which led to a final prototype based in Edinburgh. It was this prototype that was evaluated through controlled task-based experiments with real pedestrians in central Edinburgh.

Beyond this integration and evaluation work, much work was targeted at the focused areas mentioned above. 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 documents 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 has been to compare and contrast alternative techniques to solve problems within the SPACEBOOK task domain.

This project has realized a series of integrated prototypes supporting SPACEBOOK objectives culminating in a full integrated prototype that was evaluated on the streets of Edinburgh in Fall of 2013. In our deliverable D6.2.2 we present a comprehensive evaluation of this final prototype.

The project has also generated several open data sets as well as a semantic parsing corpus. Specifically we have made a public data release of our WoZ studies (linked from the project website). We shall soon publish a semantic parsing corpus for the SPACEBOOK domain based on this WoZ data. Finally the data underlying our final evaluation, as well as several side experiments, will be published on our website.

This project has generated a large body of published work [36, 8, 20, 19, 7, 1, 47, 16, 18, 17, 6, 43, 34, 21, 5, 44, 4, 14, 27, 25, 46, 9, 28, 51, 31, 29, 12, 33, 30, 26, 23, 15, 13, 32, 50, 24, 10, 11, 22, 45, 49, 42, 35, 48, 2, 3, 39, 41, 38, 37, 40] (accessible from our project website). Much of this work either documents full or partial system integrations and evaluations, or describes more directed results. To highlight, we have publications that develop and evaluate new methods of compressing belief states to enable more practical use of POMDPs for interaction management. We have publications that develop and evaluate a novel method of probabilistic behavior tracking that scales to thousands of behaviors. We have several publications that develop and evaluate methods addressing the central problem of determining which landmarks in a scene are the most salient, and thus deserving of mention to the user to assist in navigation and exploration. We have reports and publications in process that describe our approaches to advancing speech-based question answering over text sources associated with spatially-located entities. Finally, we have reports and publications in process that present and evaluate learned semantic parsers over the corpus of natural language / meaning representation language pairs we developed in this project.

In addition to our publications we have produced a series of videos that explain our project concept, showcase our city model and pedestrian tracker and demonstrate our working systems. These demonstration videos are also linked on our project website.

We expect that our findings will impact technology development for eyes-free hands-free devices for city exploration. The project has generated implementations, data, corpora and publications over a broad array

of necessary sub-components to support the SPACEBOOK vision. Partner's capabilities in these areas have been enhanced through the efforts, collaborations and outcomes developed in this project.

4 The potential impact

The SPACEBOOK project has generated concrete technical and scientific advances for eyes-free, hands-free navigation and exploration systems which 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.

Street level experiments revealed the challenges of 'dialogue based technologies'. Under favorable conditions, the concept of 'dialogue only interaction' proved to be an exciting way of exploring the city. The reliability and robustness of the system requires ASR technology that can cope with the ambient conditions of the urban world. We have developed a large body of work and expertise in this area. Various paths to market and use can be envisioned. The goal now will be for us and others to discover a way to get this technology deployed in European cities.

5 Project website

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

Project Publications

- [1] Jacob Andreas, Andreas Vlachos, and Stephen Clark. Semantic parsing as machine translation. In *ACL (2)*, pages 47–52, 2013.
(<http://spacebook-project.eu/pubs/ACL-13-1.pdf>).
- [2] Phil Bartie, Enrico Clementini, and Femke Reitsma. A qualitative model for describing the arrangement of visible cityscape objects from an egocentric viewpoint. *Computers, Environment and Urban Systems*, 2013.
(<http://spacebook-project.eu/pubs/CEUS-13.pdf>).
- [3] Phil Bartie and William Mackaness. Optimal sampling strategies for line-of-sight calculations in urban regions. In *proc. of GISRUK2012*, 2012.
(<http://spacebook-project.eu/pubs/GISRUK-12.pdf>).
- [4] Blai Bonet and Hector Geffner. Planning under partial observability by classical replanning: Theory and experiments. In *IJCAI*, pages 1936–1941, 2011.
(<http://spacebook-project.eu/pubs/IJCAI-11-B.pdf>).
- [5] Blai Bonet and Hector Geffner. Width and complexity of belief tracking in non-deterministic conformant and contingent planning. In *Proc. AAAI*, 2012.
(<http://spacebook-project.eu/pubs/AAAI-12.pdf>).

- [6] Blai Bonet and Hector Geffner. Causal belief decomposition for planning with sensing: Completeness results and practical approximation. In *IJCAI*, 2013.
(<http://spacebook-project.eu/pubs/Bonet13.pdf>).
- [7] Johan Boye, Morgan Fredriksson, Jana Götze, Joakim Gustafson, and Jurgen Königsmann. Walk this way: Spatial grounding for city exploration. In *Proc. 4th international workshop on spoken dialogue systems, IWSDS'2012*, Paris, France, November 2012.
(<http://spacebook-project.eu/pubs/IWSDS-12.pdf>).
- [8] Johan Boye, Morgan Fredriksson, Jana Götze, and Jurgen Königsmann. A demonstration of a natural-language pedestrian routing system. In *Proceedings of IWSDS'14, 5th international workshop on Spoken Dialogue Systems, Napa California.*, 2014.
(<http://spacebook-project.eu/pubs/IWSDS-14.pdf>).
- [9] Paul A. Crook, Simon Keizer, Zhuoran Wang, Wenshuo Tang, and Oliver Lemon. Real user evaluation of a POMDP spoken dialogue system using automatic belief compression. *Computer Speech & Language*, (0):–, 2014.
- [10] Paul A. Crook and Oliver Lemon. Lossless value directed compression of complex user goal states for statistical spoken dialogue systems. In *INTERSPEECH*, pages 1029–1032, 2011.
(<http://spacebook-project.eu/pubs/INTERSPEECH-11.pdf>).
- [11] Paul A. Crook, Brieuc Roblin, Hans-Wolfgang Loidl, and Oliver Lemon. Parallel computing and practical constraints when applying the standard POMDP belief update formalism to spoken dialogue management. In *International Workshop Series on Spoken Dialogue Systems Technology (IWSDS)*, 2011.
(<http://spacebook-project.eu/pubs/IWSDS-11.pdf>).
- [12] Paul A. Crook, Zhuoran Wang, Xingkun Liu, and Oliver Lemon. A statistical spoken dialogue system using complex user goals and value directed compression. In *EACL*, pages 46–50, 2012.
(<http://spacebook-project.eu/pubs/EACL-12.pdf>).
- [13] Heriberto Cuayáhuítl, Ivana Kruijff-Korbayova, and Nina Dethlefs. Hierarchical dialogue policy learning using flexible state transitions and linear function approximation. In *proc. of COLING 2012*, 2012.
(<http://spacebook-project.eu/pubs/COLING-12.pdf>).
- [14] Nina Dethlefs, Heriberto Cuayahuitl, Helen Hastie, Verena Rieser, and Oliver Lemon. Cluster-based prediction of user ratings for stylistic surface realisation. Gothenberg, Sweden, 2014. *proc of EACL*.
(<http://spacebook-project.eu/pubs/Dethlefs14.pdf>).
- [15] Nina Saskia Dethlefs, Helen Hastie, Verena Rieser, and Oliver Lemon. Optimising incremental generation for information presentation of mobile search results. In *Proceedings of the Symposium: Influencing People with Information (SIPI), 25th April, 2012, Aberdeen, UK*, 2012.
(<http://spacebook-project.eu/pubs/SIPI-12-B.pdf>).
- [16] Octavi Font, Guillem França, Anders Jonsson, Phil Bartie, and William Mackaness. Probabilistic activity recognition for navigation. In *11th Workshop on Positioning, Navigation and Communication (WPNC'14), Dresden, Germany*, 2014.
(<http://spacebook-project.eu/pubs/WPNC-14.pdf>).

- [17] Hector Geffner. Computational models of planning. *WIREs Cogn Sci*, 4:341–356, 2013. (<http://spacebook-project.eu/pubs/CogSci-13.pdf>).
- [18] Hector Geffner. Artificial intelligence: From programs to solvers. *AI Commun.*, 27(1):45–51, 2014. (<http://spacebook-project.eu/pubs/Geffner14.pdf>).
- [19] Jana Götze and Johan Boye. Deriving salience models from human route directions. In *Proc. CoSLI-3 workshop on Computational Models of Spatial Language Interpretation and Generation*, Potsdam, 2013. (<http://spacebook-project.eu/pubs/CoSLI-3.pdf>).
- [20] Jana Götze, Johan Boye, Morgan Fredriksson, and Jurgen Königsmann. Interpreting instructions in a pedestrian routing domain (extended abstract). In *Proceedings of DialDam, The 17th Workshop on the Semantics and Pragmatics of Dialogue, Amsterdam.*, 2013. (<http://spacebook-project.eu/pubs/PRAG-13.pdf>).
- [21] Patrik Haslum, Malte Helmert, and Anders Jonsson. Safe, strong, and tractable relevance analysis for planning. In *ICAPS*, 2013. (<http://spacebook-project.eu/pubs/ICAPS-13-2.pdf>).
- [22] Srinii Janarthanam and Oliver Lemon. The GRUVE challenge: Generating routes under uncertainty in virtual environments. In *Proceedings of ENLG / Generation Challenges*, 2011. (<http://spacebook-project.eu/pubs/ENLG-11-A.pdf>).
- [23] Srinii Janarthanam, Oliver Lemon, Xingkun Liu, Phil Bartie, William Mackaness, Tiphaine Dalmas, and Jana Götze. A spoken dialogue interface for pedestrian city exploration: integrating navigation, visibility, and question-answering. In *proc. of SemDial 2012*, Paris, France, September 2012. (<http://spacebook-project.eu/pubs/SEMDIAL-12.pdf>).
- [24] Srinivasan Janarthanam, Helen Wright Hastie, Oliver Lemon, and Xingkun Liu. "the day after the day after tomorrow?" a machine learning approach to adaptive temporal expression generation: training and evaluation with real users. In *SIGDIAL Conference*, pages 142–151, 2011. (<http://spacebook-project.eu/pubs/SIGDIAL-11.pdf>).
- [25] Srinivasan Janarthanam, Robin Hill, Anna Dickinson, and Morgan Fredriksson. Click or type: An analysis of wizard's interaction for future wizard interface design. In *Dialogue in Motion workshop, EACL*, 2014. (<http://spacebook-project.eu/pubs/DiM-2-14.pdf>).
- [26] Srinivasan Janarthanam and Oliver Lemon. Influencing user behaviour in personalised location based services. In *Symposium on Influencing People with Information (SIPI)*, Aberdeen, UK, April 2012. (<http://spacebook-project.eu/pubs/SIPI-12-A.pdf>).
- [27] Srinivasan Janarthanam and Oliver Lemon. Multi-threaded interaction management for dynamic spatial applications. In *Dialogue in Motion workshop, EACL*, 2014. (<http://spacebook-project.eu/pubs/DiM-1-14.pdf>).

- [28] Srinivasan Janarthanam, Oliver Lemon, Phil Bartie, Tiphaine Dalmas, Xingkun Liu, William Mackaness, and Bonnie Webber. A speech-enabled intelligent agent for pedestrian navigation and tourist information. In *Proc. IVA*, 2013.
(<http://spacebook-project.eu/pubs/IVA-13.pdf>).
- [29] Srinivasan Janarthanam, Oliver Lemon, Phil J. Bartie, Tiphaine Dalmas, Anna Dickinson, Xingkun Liu, William A. Mackaness, and Bonnie L. Webber. Evaluating a city exploration dialogue system with integrated question-answering and pedestrian navigation. In *ACL (1)*, pages 1660–1668, 2013.
(<http://spacebook-project.eu/pubs/ACL-13-2.pdf>).
- [30] Srinivasan Janarthanam, Oliver Lemon, and Xingkun Liu. A web-based evaluation framework for spatial instruction-giving systems. In *ACL (System Demonstrations)*, pages 49–54, 2012.
(<http://spacebook-project.eu/pubs/ACL-12.pdf>).
- [31] Srinivasan Janarthanam, Oliver Lemon, Xingkun Liu, Phil Bartie, William Mackaness, and Tiphaine Dalmas. A multithreaded conversational interface for pedestrian navigation and question answering. In *SIGDIAL*, 2013.
(<http://spacebook-project.eu/pubs/SIGDIAL-13.pdf>).
- [32] Srinivasan Janarthanam, Oliver Lemon, Xingkun Liu, Phil Bartie, William Mackaness, Tiphaine Dalmas, and Jana Götze. Conversational natural language interaction for place-related knowledge acquisition. In *Spatial Cognition, P-KAR workshop*, 2012.
(<http://spacebook-project.eu/pubs/P-KAR-12.pdf>).
- [33] Srinivasan Janarthanam, Oliver Lemon, Xingkun Liu, Phil J. Bartie, William A. Mackaness, Tiphaine Dalmas, and Jana Götze. Integrating location, visibility, and question-answering in a spoken dialogue system for pedestrian city exploration. In *SIGDIAL Conference*, pages 134–136, 2012.
(<http://spacebook-project.eu/pubs/SIGDIAL-12.pdf>).
- [34] Anders Jonsson, Peter Jonsson, and Tomas Lööv. When acyclicity is not enough: Limitations of the causal graph. In *ICAPS*, 2013.
(<http://spacebook-project.eu/pubs/ICAPS-13-1.pdf>).
- [35] W.A. Mackaness, P. Bartie, T. Dalmas, S. Janarthanam, O. Lemon, Liu X., and Webber B. Talk the walk and walk the talk: Design, implementation and evaluation of a spoken dialogue system for route following and city learning. In *Annual Conference of the Association of American Geographers*, 2014.
(<http://spacebook-project.eu/pubs/GEO-14.pdf>).
- [36] R Meena, J. Boye, G. Skantze, and J. Gustafson. Crowdsourcing street-level geographic information using a spoken dialogue system. In *to appear in SIGDial*, Philadelphia, USA, 2014.
- [37] Michael Minock and Johan Mollevik. Context-dependent 'near' and 'far' in spatial databases via supervaluation. *Data & Knowledge Engineering*, 86(0):295 – 305, 2013.
(<http://spacebook-project.eu/pubs/DKE-13.pdf>).
- [38] Michael Minock and Johan Mollevik. Prediction and scheduling in navigation systems. In *Proceedings of the Geographic Human-Computer Interaction (GeoHCI) workshop at CHI*, April 2013.
(<http://spacebook-project.eu/pubs/GeoHCI.pdf>).

- [39] Michael Minock, Johan Mollevik, and Mattias Åsander. Does tts-based pedestrian navigation work? Technical Report UMINF-07.14, Umeå University, 2012.
(<http://www8.cs.umu.se/research/uminf/index.cgi?year=2014&number=7>).
- [40] Michael Minock, Johan Mollevik, and Mattias Åsander. Towards an active database platform for guiding urban pedestrians. Technical Report UMINF-12.18, Umeå University, 2012.
(<https://www8.cs.umu.se/research/uminf/index.cgi?year=2012&number=18>).
- [41] Michael Minock, Johan Mollevik, Mattias Åsander, and Marcus Karlsson. A test-bed for text-to-speech-based pedestrian navigation systems. In *NLDB*, pages 396–399, 2013.
(<http://spacebook-project.eu/pubs/NLDB-13.pdf>).
- [42] Bartie P., Mackaness W.A., Petrenz P., and A. Dickinson. Automated clustering of landmark tags in urban images. In *proc. of GISRUK2014*, 2014.
(<http://spacebook-project.eu/pubs/GISRUK-14-2.pdf>).
- [43] Fabio Patrizi, Nir Lipovetzky, and Hector Geffner. Fair LTL synthesis for non-deterministic systems using strong cyclic planners. In *IJCAI*, 2013.
(<http://spacebook-project.eu/pubs/Patrizi13.pdf>).
- [44] Miquel Ramírez and Hector Geffner. Goal recognition over POMDPs: Inferring the intention of a POMDP agent. In *IJCAI*, pages 2009–2014, 2011.
(<http://spacebook-project.eu/pubs/IJCAI-11-A.pdf>).
- [45] Verena Rieser, Simon Keizer, Oliver Lemon, and Xingkun Liu. Adaptive information presentation for spoken dialogue systems: Evaluation with human subjects. In *Proceedings of ENLG*, 2011.
(<http://spacebook-project.eu/pubs/ENLG-11-B.pdf>).
- [46] Verena Rieser, Oliver Lemon, and Simon Keizer. Natural language generation as incremental planning under uncertainty: Adaptive information presentation for statistical dialogue systems. *ACM Transactions on Speech and Language Processing*, 2014.
(<http://spacebook-project.eu/pubs/Rieser14.pdf>).
- [47] Andreas Vlachos. An investigation of imitation learning algorithms for structured prediction. In *10th European Workshop on Reinforcement Learning, JMLR W&C Proceedings, Vol. 24.*, 2012.
(<http://spacebook-project.eu/pubs/EWML-12.pdf>).
- [48] Mackaness W.A., Bartie P., and C. Sanchez-Rodilla. Understanding information requirements in ‘text only’ pedestrian wayfinding systems. In *proc. of GIScience Conference Vienna*, 2014.
(<http://spacebook-project.eu/pubs/GSCI-14.pdf>).
- [49] Mackaness W.A., Bartie P., and C. Sanchez-Rodilla. The ‘when, how and what’ of text based wayfinding instructions for urban pedestrians. In *proc. of GISRUK2014*, 2014.
(<http://spacebook-project.eu/pubs/GISRUK-14-1.pdf>).
- [50] Zhuoran Wang and Oliver Lemon. A nonparametric bayesian approach to learning multimodal interaction management. In *proceedings of Spoken Language Technology (SLT)*, 2012.
(<http://spacebook-project.eu/pubs/SLT-12>).

- [51] Zhuoran Wang and Oliver Lemon. A simple and generic belief tracking mechanism for the dialog state tracking challenge: On the believability of observed information. In *Proceedings of the 14th Annual SIGdial Meeting on Discourse and Dialogue*, Metz, France, August 2013. (<http://spacebook-project.eu/pubs/Lemon-13.pdf>).