PROJECT FINAL REPORT

Grant Agreement number: FP7-287615

Project acronym: PARLANCE

Project title: Probabilistic Adaptive Real-time Learning and Natural Conversational Engine

Funding Scheme: EC FP7, Call 7

Period covered: from 1-Nov 2011 **to** 31-Oct 2014

Name of the scientific representative of the project's co-ordinator[1], Title and Organisation:

Dr Helen Hastie, Heriot-Watt University

Tel: +44 131 451 3344

Fax: +44 131 451 3327

E-mail: h.hastie@hw.ac.uk

Project website address: http://www.PARLANCE-project.eu/

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the Grant Agreement.

Executive Summary

The goal of the PARLANCE project is to design and build mobile applications that approach human performance in conversational interaction, specifically in terms of the interaction *skills* needed to do so. To realise this vision, PARLANCE has made advances in robust, incremental human-machine interaction, personalisation, and adaptation to new domains and concepts. These *skills* were mostly learned or adapted using real data, and used to build systems for voice-enabled interactive hyper-local search and information provision applications.

The project was in two stages where for System 1, we created an incremental architecture that uses a micro-turn Interaction Manager (IM) to create more fluid interactions. We also looked at statistical approaches to the problem of incremental Natural Language Generation (NLG) and Text-to-speech (TTS). System 1.5 built upon this work further by adding in an improved Spoken Language Understanding (SLU) and IM policy. We were able to show good progress with System 1.5 having greater understanding and shorter dialogues than System 1 as well as significantly higher task completion over the baseline using a learned policy.

The goal of System 2 was to develop methods for automatically adapting to extended domains. This work will enable systems to grow and adapt automatically to new concepts rather than having to be re-engineered each time. This is very important, particularly for information provision applications where the domain is constantly evolving and changing. In order to capture this change, we developed dynamic, modular ontologies as part of the knowledge base. As well as component based evaluations, we evaluated the domain adaptation methods developed on the project by incorporating them into an end-to-end system and evaluating against a TopBound system trained on data from the extended domain. Here, we met our objectives by showing no significant differences for objective Task Completion as well as no significant differences perceived for complex tasks in terms of answers to four subjective questions. Finally, the adaptive systems actually have shorter dialogues than the TopBound system. A second goal of System 2 was to investigate how to adapt dialogue to the user. In this regard, we developed a schema for a User Model and showed that it is possible to effectively personalise both surface realisation and interaction policy to a user with good results, as shown in component evaluations.

In terms of multilinguality, we produced PARLANCE apps in English, Mandarin and Spanish and demonstrated these at conferences with excellent feedback. Finally, in order to further investigate spoken dialogue systems and in particular personalisation techniques, we believe that the best forum for this is evaluation “In the Wild” with real users with real preferences doing real tasks. To this end, we created a framework for evaluating spoken dialogue systems through an App available on the Google Play store. This App collects some information on the user (e.g. age and preferred cuisine) and uses Yahoo’s Hyperlocal API to enhance the presentation of recommendations and to provide result-set specific ordering and filtering options. We believe this framework is a first of its kind in the scientific community enabling incremental improvements for research systems with real users.

Finally, a large repository of data were collected over the course of the project totaling approximately 130,000 utterances annotated for various aspects of dialogue including dialogue acts/semantics, transcription and turn-taking features. This is a valuable resource and will be made available to the community

Project Concept and Objectives

The project goal is to design and build mobile applications that approach human performance in conversational interaction, specifically in terms of the interactional skills needed to do so, such as recognising and generating conversational speech incrementally in real-time, adapting to new information and reacting to new communicative situations. All of these skills that we have developed were learned or adapted using real data, and used to build systems for voice-enabled interactive hyper-local search and information provision applications. Current search engines work well only if the user has a single search goal and does not have multiple trade-offs to explore. For example, standard search works well if you want to know the phone number of a specific business but poorly if you are looking for a house with several different search criteria of varying importance, e.g. number of bedrooms versus bathrooms versus price etc. The latter requires the user to collaborate conversationally over several turns.

Overall Project Objectives

There are four main objectives of this work that aim to revolutionise the way people search for information and the way spoken dialogue systems are developed and maintained:

- O1: develop incremental, responsive dialogue systems in 3 languages
- O2: develop personalised dialogue systems that adapt to different users with different goals in different contexts,
- O3: develop dialogue systems that are dynamic and evolve towards open-ended domains,
- O4: develop interactive hyper-local search.

O1: Develop incremental, responsive dialogue systems in 3 languages
For spoken dialogue systems to be universally adopted, advances in the state-of-the-art are necessary to enable highly responsive and conversational systems. Traditionally, the unit of speech has been an utterance with strict, rigid turn-taking. In contrast, incremental dialogue systems that deal with micro-turns will be more fluid and responsive. Instead of treating user barge-in as a special case to be handled by the signal processing and otherwise hidden from the Interaction Manager, user barge-in will be treated as a normal speech act and an intrinsic part of everyday collaborative conversational problem solving. To complement this highly responsive interaction, there is a need for real-time, accurate analysis of natural language queries.

O2: Develop personalised dialogue systems that adapt to different users with different goals in different contexts
For a system to be truly natural it should adapt its interaction style to whom it is speaking, just as humans do. To this end, PARLANCE objectives include designing and building a User Model and investigating personalisation methods for the various components including the IM and NLG.

O3: Develop dialogue systems that are dynamic and evolve towards open-ended domains
Currently, spoken dialogue systems are created for a specific domain using fixed databases (for example, bus schedules, restaurants) making them rigid and making new domains costly to develop. The project objective was to develop algorithms to enable spoken dialogue systems to incorporate new concepts associated with open-ended domains and evolve rather than be developed for specific applications. The term “open-ended domains” refers to the system’s ability to handle the introduction of new words and concepts in a single domain and in addition facilitate easy transfer from one domain to another.

O4: Interactive, hyper-local search
The project aims to create the first truly interactive search system that provides search results that are personalised, social, and location-aware. This can be broken down into three parts. Firstly, structured information has to be extracted from the web to populate dynamic ontologies and match this structured information to the user's natural language query. Secondly, mobile, hyper-local information will be used to instruct both the search and the interaction.

Objectives in Terms of End-to-end Systems

During the first part of the project, the goal for PARLANCE System 1/1.5, was to create a system that is incremental (01). Traditionally, the unit of speech has been a whole utterance with strict, rigid turn-taking determined by a voice-activity detector. By creating a system that is incremental, we hoped to improve the experience for the user and create more natural interactions.

Objectives addressed in the second part of the project, included techniques for personalisation (O2), automatically adapting to new domains (O3) and show casing these in apps for 3 different languages that leverage hyper-local search (O4/O1). The project objectives include learning of interactional skills for coping with unforeseen concepts thus moving beyond the current focus on limited application domains that use fixed, static databases. In addition, an objective was to develop dynamic User Models that enable the system to adapt to the user thus personalising interaction.

Finally, project objectives include running component evaluations throughout the 3 year project to evaluate individual techniques for individual modules with respect to the four objectives listed above. In addition, a number of end-to-end system evaluations were conducted including 3 variants of System 1 for incrementality and also 3 variants of System 2 for adapting to new domains in the information provision domain and hyperlocal search, specifically for searching for restaurants.

As well as these original objectives, an additional task was added at month 25 to develop and App for the public to use to facilitate evaluation "In the Wild".

Objectives by Work Package

Overall Objectives of WP1:

- develop a micro-turn architecture, incorporating methods for both producing and interpreting speech recognition hypotheses incrementally during the course of an utterance;
- develop ASR and SLU modules in this architecture that support natural conversational speech in multiple languages;
- incorporate dynamically changing vocabulary based on information in the User Model (e.g., user preferences, context, and geographic location) and search results.

Overall Objectives of WP2:

- support decision making at a micro-turn level in order to provide natural, conversational interaction;
- adapt to a dynamic ontology of the domain of discourse and a dynamic knowledge base (developed in WP4) in order to gracefully handle new topics and searches; and
- adapt to different users in real-time, in order to optimise the dialogue strategy for that user.

Overall Objectives of WP3:

- develop highly responsive output components through incremental output planning and execution;
- optimise language generation and synthesis for individual users, thus providing a personalised experience; and
- adapt to dynamic ontologies and be able to portray constantly changing information to the user in a clear and concise manner.

Overall Objectives of WP4:

- incrementally build a dynamic, modular "ontological" knowledge base; and
- build and maintain a constantly changing User Model and in particular look at social preferences and context.

Overall Objectives of WP5:

- provide the back-end search services that support an interactive hyper-local search experience;
- discover local content in an automated way from alternative sources, such as open gazetteers, user-generated content (such as Flickr and Twitter), and Web content;
- refine the search by refining the content that is available to serve to the user, and by understanding the user needs by examining his/her personal preferences and social context;
- understand the user's geographic context to determine the scope and intent of the information they seek,
- provide an API, or set of services, that can be used to build a hyper-local search application; and
- develop a multimodal, mobile application (additional objective from DoW).

Overall Objectives of WP6:

- provide requirements analysis, architecture design and integration for Systems 1 and 2,
- conduct evaluations and data collections for Systems 1 and 2; and
- develop a multimodal, mobile application, modify the PARLANCE server to run multiple instances and hook up to the PARLANCE App available on the store (additional objective from DoW).

Main S&T Results/Foregrounds

Here we discuss the main PARLANCE results in terms of System 1 which addresses the first objective of incremental dialogue in three languages and then of System 2 which looks at the other three objectives of personalisation, adaptivity to new domains and hyperlocal search. In the following section, we detail the progress over the 3 years of the project. We then provide further detail on the results by work package, give a description of the data repository and then summarise the final results.

Summary of results by System

System 1: Incremental

The impetus during the first year and a half of the project was on the development and integration of core System 1 components to provide an end-to-end incremental system as reported in D6.3 (O1). This system incorporates all the modules as described in the architecture reported in D6.1. The resulting system has a unique micro-turn architecture as reported in D1.1 (see Figure 1). This micro-turn speech subsystem consists of a micro-turn controller that uses the incremental ASR output to decide whether to signal a turn change or not. The user turn is signalled when a non-filler word is found in the incremental output. While in user turn state, if the incremental output offers limited information, the system outputs a short back-channel and stays in user turn state. This involved developing fundamental software components necessary to support incremental ASR by performing two passes, SLU and micro-turn Interaction Management, as well as a separate Voice Activity Detection module (see D1.1 for details).

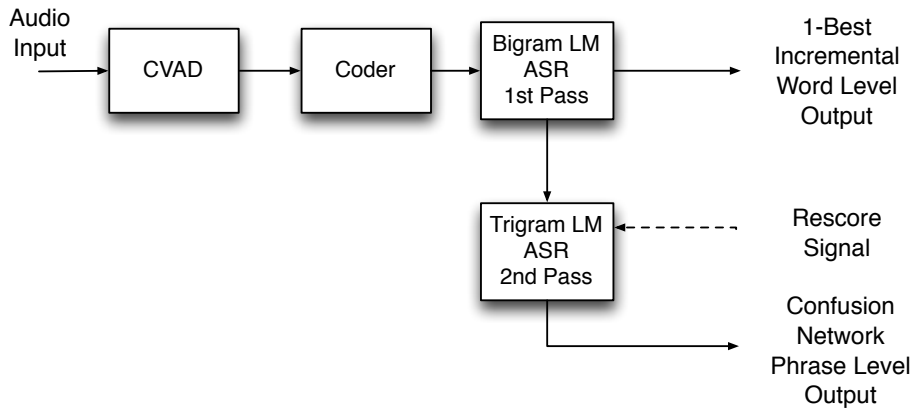

Figure 1. Micro-turn ASR Configuration

The PARLANCE INCremental Dialogue Act (PINC) scheme (see D6.1) was developed, which decomposes traditional dialogue acts into smaller units or *primitive* dialogue acts that can be recognised incrementally from partial user utterances.

The NLG component integrated into System 1 incorporates a state-of-the-art surface realiser that uses Conditional Random Fields (D3.1). This prototype extends the state-of-the-art in several ways, most significantly through the application of statistical techniques for the automatic optimisation of system behaviour and by being more context-aware and human-oriented than current incremental surface realisers. A crowd-sourcing evaluation showed this method to have more positive feedback when compared against a number of baselines.

The architecture of the Knowledge Base including the Ontology Manager and User Model were defined and a novel approach to ontology module population was developed using dependency analysis and a web search engine (D4.2/D4.3). Pattern detection techniques for text-based information retrieval were developed and extensive experiments were carried out to enrich modular ontologies automatically with new attributes and values by mining the Web. Using this method, the accuracy of the discovered attribute values for 15 attribute types is between 72-87%, using a corpus including 50 snippets for each attribute type. Finally, an improved method for bootstrapping relations from a small set of seed tuples was developed, which potentially outperforms state-of-the-art iterative bootstrapping methods, scales well to large corpora, and requires only minimal linguistic analysis.

As well as creating a search service for the domain as an interactive hyper-local search API which was done ahead of schedule (D5.2), there have been key developments in hyperlocal search including methods for discovering points of interest in unstructured text; assigning geographic coordinates to a point of interest within a median accuracy of 0.96 kilometres; and categorising points of interest with 70% accuracy (D5.1).

System 1 was evaluated using Crowdflower crowdsourcing platform where 1,774 dialogues were collected and annotated. System 1.5 was an improved version of System 1 evaluated in the third year of the project. The goal of evaluating System 1.5 was to perform a short evaluation of the system with the improved SLU, a new IM learned policy and also to pilot the new Crowdflower set-up. In this configuration, we used Gold Standard questions whereby the users had to answer the code correctly, if not then they were kicked out of the experiment. This was proven to eliminate the problem that users were cheating or hacking into the system. Results of System 1.5 evaluation showed how good progress was made, with System 1.5 having greater understanding and shorter dialogues than System 1 (16.4 average turns compared to 20.5) as well as significantly higher task completion over the baseline for System 1.5A that uses a learned policy. Subjective task success was at highest 89% (see D6.7 for details).

Finally, during the first part of the project good progress was made towards a Mandarin version, which involved translating NLG and SLU component and back-end datasets.

System 2: Adaptation to new domains, personalisation and hyperlocal

Work towards System 2 objectives (O2, O3 and O4) continued into Year 3 and excellent progress was made, particularly with respect to adapting to new domains. With the advice of expert reviewers the project did shift emphasis more towards mobile applications and collecting real data. PARLANCE Apps had been developed for English, Mandarin and Spanish. However, these apps used third party Google TTS/ASR and did not allow us to evaluate the full functionality of the PARLANCE system and, in addition, did not allow for recording of the dialogues and proper logging. A framework was, therefore, developed for English that used audio compression to communicate with the PARLANCE server. The intention is that the App will be put on the Google Store for the general public to use (See Figure 2 for screenshots). Much effort was put into making this App robust enough for public use and allowing multiple instances of the system to run. In addition, hyper-local features were integrated into the mobile application and were used to optimise search results (O4).

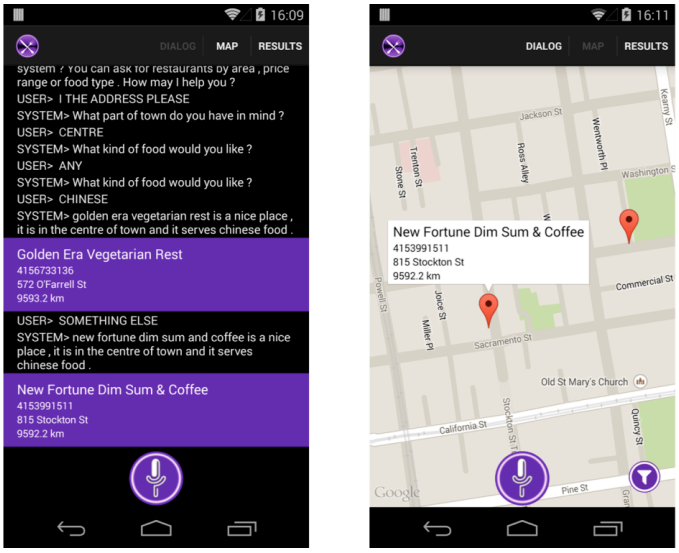

Figure 2: Final version of the English app

With regards adaptive ASR and SLU, we exploited neural network models trained on large open-domain datasets to improve the ability of models trained on one domain to adapt well to a new domain (O3). For ASR, recurrent neural networks trained as large-vocabulary language models were combined with training on a narrow domain to improve robustness when applied to related but new domains (D1.3). For SLU, we exploited neural network word representations (called embeddings) trained on large open-domain datasets inside our neural network parsing model trained on a narrow domain. These word embeddings helped SLU understand terminology from a new domain that was not in the training data. A redesign of the SLU model even allowed it to output novel semantic terms by using word embeddings to represent these terms (D1.3). These innovations for adaptability were combined with advances in the basic ASR and SLU models that also improved accuracy and speed (D1.2).

The Interaction Manager was expanded to quickly learn how to talk about the new concepts using its knowledge of the existing concepts (O3). The problem is two-fold, the Interaction Manager must be able to accurately maintain belief over new concepts and adapt its policy to be able to maintain the conversation about the new concepts. Three ontologies of increasing complexity were generated in collaboration with WP4 and we adopted an incremental scheme, which enables the domain of the Interaction Manager to be repeatedly extended by recursively specifying priors in Gaussian processes. We showed that it is possible to effectively double the number of concepts or attribute types understood by a system providing restaurant information using only 1,000 adaptation dialogues with real users (see D2.2 for further details). This was shown to have no significant difference in performance compared to a TopBound system trained for the expanded domain for complex dialogues for objective Task Completion (D6.7).

With regards the system being able to talk about new concepts, a statistical surface realiser was trained from unlabelled data (utterances in raw text) through automatic semantic annotation for easy domain adaptation. A subjective evaluation reported that training it from unlabelled data can generate as good results as training it from labelled data. In the end-to-end evaluation this was also shown to have no significant differences compared to the TopBound system for objective Task Completion (D6.7). Regarding speech synthesis, an in-domain corpus was developed for expressive speech synthesis targeted to the application, and an expressive HMM voice was trained (D3.4).

Finally, we conducted a survey of the state of the art with respect to conversational question-answering systems and ontologies for spoken dialogue systems, together with an overview of useful linguistic and semantic tools (D4.4). In addition to this work, we provided a method for dealing with vagueness in relation to question-answering systems.

In terms of personalization (O2), we developed a schema for a User Model and extended it to handle sentiment analysis for a finer-grained detection of user preferences and provided a User Model similarity metric and a framework for integration of vague concepts. In terms of personalizing interaction, the Interaction Manager policy trained on a simulated user was shown to be improved via interaction with real users showing the capability of the system to adapt to different users. This personalised Interaction Manager significantly outperforms the standard manager when tested with users with clear preferences in simulation. With regards personalised NLG for System 2, we showed that we can entrain the system to reflect stylistic features of the user. Specifically, a statistical surface realiser was trained from stylistic sentences and user ratings in three dimensions (naturalness, colloquialism, and naturalness). A subjective evaluation reported that the generated styles are recognisable to human judges and therefore would be useful in an end-to-end system.

In terms of end-to-end evaluations for System 2, we evaluated 3 variants of system 2 enabling comparison of IM and NLG modules (see D6.7). These variants were compared against a TopBound system trained on data on the extended domain. Here, we met our objectives by showing no significant differences for objective Task Completion as well as no significant differences perceived for complex tasks in terms of answers to four subjective questions. Subjective Task Completion was between 88-95%. Automatically calculated Task Completion was a lot lower than this, however, it has been shown that this metric should really be hand-labelled as subjects do have a tendency to change their goals which is not captured by the Task Completion script (see D6.4). Nevertheless this metric can be used to compare systems. Finally, the adaptive systems actually have shorter dialogues than the TopBound system. In terms of evaluation of personalisation, as mentioned above, we showed it is possible to effectively personalise both surface realisation and interaction policy to a user with good results. However, emphasis was shifted more towards an App that can be used to evaluate aspects of personalisation “In the Wild” with real users doing real tasks.

Summary of Results by Year

Here we summarise work done by year and then also by work package.

Work performed in Year 1:

- System Architecture design and requirements analysis (D6.1);
- Evaluation strategy and metrics for System 1 (D6.2);
- Micro-turn architecture to enable incremental speech recognition and natural language understanding: design of architecture and initial prototype (D1.1);
- New statistical approaches to incremental Spoken Language Understanding (D1.1);
- Initial architecture of the adaptive Knowledge Base component, including the Ontology Manager and the User Model (D6.1); and
- Geographically aware content discovery (D5.1).

All deliverables and milestones were reached: D1.1, D5.1, D6.1, D6.2, D7.1, D7.2, D8.1, D8.2.1, D8.2.2, D8.3.1, D8.3.2 and MS1, MS2.

Work performed in Year 2:

System 1

- Integration, test and evaluation of end-to-end system (D6.3/D6.4);
- Development of an End-Of-Turn classifier to augment the Voice Activity Detection (VAD) module for improved accuracy and latency for detecting the end of the user's turn;
- Further development, implementation, and evaluation of ASR and SLU modules for incremental interaction;
- Development of a POMDP Interaction Manager (IM) and simulator to handle micro-turn dialogues and learn appropriate strategies (D2.1);
- Development of new statistical approaches to incremental Natural Language Generation and TTS (D3.1);
- Development of an interactive hyper-local search API (D5.2) and a version of System 1 that uses the Yahoo API directly from the IM; and
- Completions of initial version of System 1 for Mandarin.

System 2

- Development of initial methods for the IM to adapt to new domains;
- Development of an initial version of Personalised NLG;
- Development of methods for open information extraction that combine shallow syntactic and semantic features (D4.1) and incremental construction of modular domain ontologies for spoken dialogue systems in open-ended domains (D4.2);
- Development of an initial version of the User Model (D4.3); and
- Development of an evaluation strategy for System 2 (D6.4).

All deliverables and milestones were reached: D2.1, D3.1, D4.1, D4.2, D4.3, D5.2, D6.3, D6.4, D6.5, D7.3, D8.2.3, D8.2.4, D8.3.3 and MS3, MS4, MS5.

Work performed in Year 3:

System 1.5

- Development and integration of a syntactic parsing based SLU module that is 40 times faster with equivalent accuracy, along with a new IM policy and more secure Crowdflower evaluation platform; and
- Evaluation showing that there is a significant improvement using a learned dialogue policy vs a hand-crafted baseline (D6.7) and that understanding was improved.

System 2

- Development of neural network models for language modelling for ASR and parsing for SLU that improve accuracy, speed, and portability to new domains (D1.2 and D1.3);
- Development of methods for the IM to adapt to new domains and personalisation (D2.2, D2.3 and D6.7);
- Development a method of personalising surface realisation (D3.2) and a method of novel and domain-transferable methods for automatic semantic slot labelling for NLG (D3.3);
- Development of adaptive HMM-based speech synthesizer (D3.4) and domain-transferable approach to training synthetic voices for target domains (D3.3);
- conducted a survey of state of the art with respect to conversational question-answering systems and ontologies for spoken dialogue systems, together with an overview of useful linguistic and semantic tools (D4.4);
- Updated the User Model to include information from sentiment analysis (D5.3);
- Investigation of interactive social search and collaborative filtering (D5.3);
- Performed end-to-end evaluation of System 2 showing that the PARLANCE techniques for domain adaptation can perform as well as a carefully crafted system trained on target domain data (D6.6/D6.7); and
- Developed 3 versions of the PARLANCE App – Mandarin and Spanish for demonstration and dissemination and an English version to be deployed in the store that works of 3G and Wi-Fi (D5.4, D6.7).

All deliverables and milestones were reached: D1.2, D1.3, D2.2, D2.3, D3.2, D3.3, D3.4, D4.4, D5.3, D5.4, D6.6, D6.7, D7.4, D8.2.5, D8.2.6, D8.2.7 and MS6, MS7, MS8.

Summary of Results by Work Package

WP1 concerns itself with ASR and SLU. In summary, WP1 contributed the following significant outcomes during the project:

- developed fundamental software components necessary to support ASR and SLU in the micro-turn architecture;
- developed an End-Of-Turn classifier to augment the Voice Activity Detection module for improved accuracy and latency for detecting the end of the user's turn;
- developed an initial prototype of the complete micro-turn ASR-SLU subsystem;
- developed the PINC dialogue act scheme, designed to allow incrementally-recognisable dialogue act components to be output by SLU to the Interaction Manager;
- evaluated the micro-turn architecture in end-to-end evaluation (System 1/System 1.5);
- developed continuous ASR module and improved ASR with updated acoustic and language models for English;
- developed initial version of Mandarin acoustic and language models;
- translated template based SLU for Mandarin and Spanish;
- obtained success in using recurrent neural networks for language models, and adapting them to new domains;
- developed syntactic parsing based SLU module that is 40 times faster with equivalent accuracy (System 1.5);
- production of a relevant corpus with high quality annotation for the SLU task, which will be made freely available; and
- successful exploitation of unsupervised word embeddings and representation learning models to improve the accuracy and portability of SLU, combined with supervised training.

WP2 concerns itself with Interaction Management (IM). In summary, WP2 contributed the following significant outcomes during the project:

- development and evaluation of a statistical micro-turn manager to support natural interaction including backchannels, barge-in by either party (System 1);
- completed a preliminary investigation into the use of inverse reinforcement learning for turn management;
- developed an incremental user simulator aimed for incremental dialogue policy learning;
- improved classification accuracy in dialogue act recognition involving barge-in using partial context comparing to using full context;
- created simulated dialogues with barge-in which are more realistic than simulated dialogues without barge-in events;
- extended Bayesian network parameters and dialogue polices to deal with dynamic ontologies;
- developed and evaluated an Interaction Manager that is able to adapt on-line in direct interaction with real users to changes in the dialogue domain (System 2);
- developed an Interaction Manager policy trained on simulated users which can further be improved via interaction with real users showing the capability of the system to adapt to different users;
- developed fast incremental on-line reinforcement algorithms which can be used to learn Interaction Manager policies either from scratch or using an existing policy as a prior; and
- developed a personalised Interaction Manager that significantly outperforms the standard manager when tested with users with clear preferences in simulation.

WP3 concerns itself with Natural Language Generation (NLG) and Text-to-Speech Synthesis (TTS). In summary, WP3 contributed the following significant outcomes during the project:

- developed incremental NLG through the application of reinforcement learning in flat and hierarchical settings;
- established a connection between the psycholinguistic principle of information density and overlaps in spoken dialogue;
- applied undirected graphical models to incremental surface realisation to ensure coherence across micro-turns and the inclusion of incremental discourse phenomena;
- developed a large corpus for the training of the HMM-based synthesiser;
- developed a prototype TTS system for Mandarin;
- developed template based NLG for Spanish and Mandarin systems
- developed a trainable NLG component that treats surface realisation as a sequence labelling task based on conditional random fields; evaluated this component with human subjects and showed that it outperforms the baseline in terms of objective task success in on-task dialogues (System 1);
- developed a grammar induction algorithm that learns new generation spaces from minimally labelled data and can be easily ported to new data sets (demonstrated on three data sets) (System 2);
- extended text analysis to include dialogue act semantic information and an emphasis assignment algorithm;
- trained dialogue context sensitive HMM-Based synthetic voice;
- designed and carried out a preference listening test to compare synthetic voices in the context of dialogue systems add more relating to TTS;
- developed a novel and domain-transferable method for estimating and modelling individual user preferences in terms of surface realisation, including stylistic dimensions colloquialism, politeness and naturalness as well as qualitative dimensions phrasing and grammaticality;
- conducted experiments alternating the speaking rate of the TTS and showing that (a) users show an overall preference for the normal speaking rate, and (b) they show alignment towards the system's speaking rate;
- developed a novel and domain-transferable method for automatic semantic slot labelling. Experiments have compared different similarity metrics as well as different amounts of supervision. In an experiment with human judges, a surface realiser trained from automatic labels received comparable ratings to one trained from human labels; and
- developed a novel and domain-transferable approach to training synthetic voices for target domains when predominantly data from a distinct source domain is available. Results have shown that while the use of Factorised Decision Trees has no impact on the source domain, user perception in the target domain is significantly improved.

WP4 concerns itself with dynamic ontologies for natural spoken interaction in open-ended domains. In summary, WP4 contributed the following significant outcomes during the project:

- defined the architecture of the Knowledge Base including the Ontology Manager and User Model. Dependencies between the components of PARLANCE and their interfaces were specified;
- designed a modular Tourism ontology connecting modules;
- designed a meta-schema for ontology modules;
- developed an approach of ontology module population using dependency analysis and a web search engine;

- developed specifications for the representation of the user (i.e., the meta-model) and ontology data in the form of XSD schemas;
- carried out extensive experiments to enrich modular ontologies automatically with new attributes and values by mining the Web. This involves pattern detection techniques for text-based information retrieval. The accuracy of the discovered attribute values for 15 attribute types is between 72-87%, using a corpora including 50 snippets for each attribute type;
- developed an improved method for bootstrapping relations from a small set of seed tuples, which potentially outperforms state-of-the-art iterative bootstrapping methods, scales well to large corpora, and requires only minimal linguistic analysis;
- implemented and integrated the modular ontologies and enrichment approach in the Knowledge Base;
- defined the policy for deriving user preferences from the dialogue history and proposed a formal concept analysis-based approach for deriving association rules for social context;
- incorporated a notion of information vagueness in the User Model;
- proposed a critique-based query refinement mechanism with adjusted user preferences;
- conducted a survey on case-based and rule-based reasoning for natural language query analysis;
- conducted a survey of state of the art with respect to conversational question-answering systems and ontologies for spoken dialogue systems, together with an overview of useful linguistic and semantic tools;
- developed a methodology to map existing domains to new domains in order to adapt PARLANCE to new domains;
- developed a tool for generating data for these new domains that can be used for supervised and unsupervised learning techniques for NLG and SLU components; and
- developed a method for dealing with vagueness in relation to question-answering systems.

WP5 concerns itself with interactive hyperlocal, social search for Spoken Dialogue Systems. In summary, WP5 contributed the following significant outcomes during the project:

- developed a method for discovering points of interest in unstructured text;
- developed a method for assigning geographic coordinates to a point of interest within a median accuracy of 0.96 kilometres;
- developed a method for categorising points of interest with 70% accuracy;
- developed a curated database of points of interest from licensed data;
- developed a curated database of points of interest from sponsored listings;
- developed a search service dedicated to manually curated restaurants;
- developed a set of search functionalities enabling users to retrieve restaurants matching their specific needs;
- developed a method of discovering local content in an automated way from alternative sources, such as open gazetteers, user-generated content (such as Flickr and Twitter), and Web content;
- refined the search process by refining the content that is available to serve to the user, and by understanding the user needs by examining his/her social context;
- developed a method of understanding the user's context to determine the scope and intent of the information they seek; and
- provided an API, or set of services, that can be used to build a hyper-local search application;
- developed the PARLANCE mobile App with an attractive interface showcasing hyper-local search features;

- developed a revised client/server architecture where compressed audio data is exchanged between the client and the PARLANCE server that can support multiple concurrent users;
- provided an overview of user similarity metrics, built upon the social search method;
- augmented the user model interest inference by adding a sentiment analysis approach;
- provided a definition of user vagueness models and how they can be used in relation to collaborative filtering; and
- conducted a pilot study into re-ranking identified instances (i.e. restaurants) in real-time by querying and analysing reviews from the web.

WP6 concerns itself with Platform Integration and Component Integration. In summary, WP6 contributed the following significant outcomes during the project:

- provided analysis and definition of user and system requirements;
- provided initial design of the PARLANCE system architecture for Systems 1 and 2;
- defined the evaluation strategy and plan for PARLANCE System 1 and System 2 well ahead of the actual evaluation;
- set up an SVN for integration and collaboration;
- developed the reference implementation of the architecture for System 1, and integration of the different components: SLU, ASR, IM, and KB;
- set up the crowdsourcing testing environment of the PARLANCE System;
- conducted evaluation and analysis of results for PARLANCE System 1, 1.5 and System 2;
- created a rich data repository of dialogues that is available to the scientific community;
- integrated, tested and evaluated three versions of System 2 for investigating adaptivity to new domains;
- developed App versions in Mandarin, English, Spanish using third party ASR/TTS; and
- conducted additional work on an Android mobile application for the store in English for San Francisco including:
    - A new communication protocol for interacting with the PARLANCE server over low-bandwidth network connections;
    - Server that can cope with multiple instances of the system to cope with traffic; and
    - Multi-modal features that allow interaction with the results using touch commands without interrupting the conversation.

Overview of Data Repository

Access to the data repository is by request though this website –linked from main PARLANCE website: https://sites.google.com/site/parlanceprojectofficial/home/datarepository

Dataset		**System**	**Num total Utterances (System and User)**	**Attributes (# values**	**Hand-annotation**	**Notes**
SFCore		1	~4,800	*food*(59), *area*(155), *pricerange*(3)	Word transcription, task success, Dialogue Acts/ Semantics,DB translated for Mandarin and Spanish	End-to-End Eval
SFCore1.5		1.5	~10,000	*food*(59), *area*(155), *pricerange*(3)	None	End-to-End Eval
SFExt	SF1Ext	2	~8,000	SFCore+ *near(39)*	Dialogue Acts/ Semantics	Data used for algorithm dev of a number of modules
	SF2Ext	2	~6,500	SF1Ext + *allowedforkids (2)*	Dialogue Acts/ Semantics	
	SF3Ext	2	~6,300	SF2Ext + *goodformeal(4)*	Dialogue Acts/ Semantics	
SFExtEval		2	~15,000	as SFExt	Word transcription	End-to-End Eval
Spacebook		N/A	~1,700	N/A	Dialogue Acts/ Semantics, turn-taking features	Leveraged for Task 3.1
SFExt_generic		2	~9,000	as SFExt	None	Data used for algorithm dev for Task 2.2
SFExt_generic1		2	~8,000	as SFExt	None	
SFExt_generic2		2	~8,200	as SFExt	None	
SFExt_scratch		2	~9,200	as SFExt	None	
SFExt_scratch1		2	~9,500	as SFExt	None	
SFExt_scratch2		2	~10,000	as SFExt	None	

Summary of Final Results

In summary, PARLANCE has made progress in the field in modules ASR, SLU, IM, NLG, TTS, search and dynamic modular ontologies. In terms of incremental dialogues, we were able to develop a system that flowed more naturally and allowed for barge-ins. One barrier to wide spread adoption of Spoken Dialogue Systems is their domain specificity and the steep curve to transfer to new domains. We successfully addressed this issue by developing algorithms that allow us to boot strap onto existing data or systems with comparable performance to carefully crafted systems for the extended domain. With regards personalisation, we have shown through component development and evaluation much potential for systems to be able to adapt to the person they are interacting with.

Finally, we have developed a suite of apps in three languages that take advantage Yahoo's API for hyperlocal search through dialogue and multimodal interaction. We have developed a number of integrated end-to-end prototypes, which have been evaluated and compared to competitive baselines with excellent results. The dialogues from these evaluations have fed into a rich data repository, which is available to the community. The consortium feels that we have met the objectives of the project with no deviations and we look forward to contributing further to the field through the app, which will remain on the store collecting further data.

Ethical Consideration and Informed Consent

The evaluations described here do not involve working with vulnerable subjects or children. There is very little risk to subjects as they are only using a phone and keyboard/mouse. With regards to consent, crowdsourced workers sign up to the terms and conditions of their chosen crowdsourcing platform (for an example see https://www.mturk.com/mturk/conditionsofuse for Amazon Turk and http://www.crowdflower.com/legal for Crowdflower). Additional information on data collection and privacy are given to the workers (see Appendix A in D6.4 for example text) including that all data will be anonymised, giving them the option to hide their phone number and informing them that this number will not be used to call them back. This experimental procedure was passed by the Heriot-Watt University Ethics committee.

Potential Impact and Use

PARLANCE technology has the potential to become embedded in everyday life for millions of people. As the amount of information on the web increases at a huge rate, in the near future, interactive, personalised search will be a key medium to digest and explain this information to the overwhelmed user. Natural, interactive search will provide vital conversational information access to many sectors of the population.

The PARLANCE programme targeted significant impact in the scientific community and in commercial applications of speech technology. Through its work on incrementality, data-driven techniques and adaptivity, it will enhance the naturalness of spoken dialogue systems, making them more likely to be adopted in everyday life. In addition, PARLANCE moves away from the traditional method of creating systems, reducing development costs and time-to-market, and stimulating innovation and expanding markets. PARLANCE is also expected to have a significant impact on the European competitive position in a multilingual digital market by providing improved services to citizens.

Main Dissemination Activities

Consortium members have undertaken the following dissemination activities, which can be broken down into the three target audiences as follows:

Academic Community	57 Peer-reviewed Publications	30 Invited Talks
	7 Workshops Organised	10 Keynote Talks
	1 Conference Organised	2 Inaugural lectures
	2 Guest Lectures	1 Tutorial
	5 Conference Sponsorships	2 EU scientific communities joined
Industry	8 Industry Visits	8 Invited Industry Talks
	1 Keynote Talk	1 Presentation
	Flyer	Fact Sheet & Slides
Wider Public	3 Press Releases	2 News Articles
	Radio Programme	TV interview
	Website	Twitter
	Facebook	LinkedIn

Academic Community

Publications

One metric for PARLANCE dissemination is the number of peer-reviewed publications for conferences and journals. These papers go through a very strict review process by experts in the field, with some conferences and journals having acceptance rates of 20% and below, indicating that only material of the highest quality is published.

PARLANCE has been published in 54 peer-reviewed conferences, workshops & journals (including winning the Best Paper Award at SIGDIAL 2013), 2 book chapters and 1 technical report. For all conferences where work was published in the proceedings, a PARLANCE consortium member gave a talk or poster presentation thus disseminating around the globe. The full list of publications can be found on the PARLANCE website with links to the papers and in Section A of this report.

Talks

Over the course of the project, members of the PARLANCE project consortium have given 30 invited talks and 10 keynote speeches in 12 different countries round the world. A full list of talks given can be found in Section B of this report. Some highlights include:

Steve Young at EUSIPCO 2013	Marrakech, Morocco	
Peter Mika at ISWC 2013	Sydney, Australia	
Peter Mika at WWW 2013	Rio de Janerio, Brazil	
Heriberto Cuayahuitl at AINL 2014	Moscow, Russia	
Steve Young at SIGDIAL 2014	Philadelphia, USA	
Marie-Aude Aufaure at ER 2013	Hong Kong	

Heriberto Cuayahuitl giving his Keynote Talk at AINL in Moscow (January 2014)

Peter Mika giving his Keynote Talk at 3rd Russian Semantic Web Conference in St Petersburg, Russia (October 2012)

Helen Hastie presenting a poster at SemDial in Edinburgh, UK (September 2014)

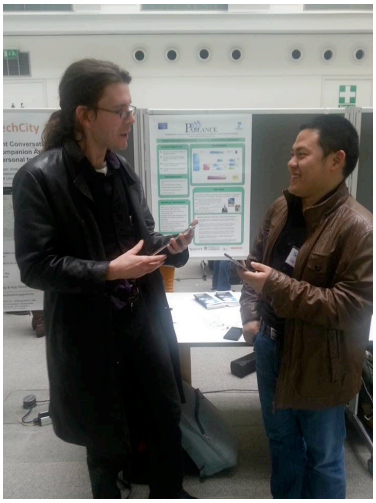

Yanchao Yu presenting a poster at SICSA Speech Technology Day in Edinburgh, UK (March 2014)

Sponsorship

PARLANCE has sponsored the following conferences which meant having presence in a variety of forms including on the conference website, on leaflets, in proceedings, programmes and having the logo on conference bags:

MLIS 2013	http://mlis-workshop.org/2013/Welcome.html	Beijing, China August 2013
MLIS 2014	http://mlis-workshop.org/2014/Welcome.html	Quebec, Canada July 2014
SemDial 2014	http://www.macs.hw.ac.uk/InteractionLab/Semdial/index.html	Edinburgh, UK September 2014
SIGDIAL 2014	http://www.sigdial.org/workshops/conference15	Philadelphia, USA June 2014
ISWC 2014	http://iswc2014.semanticweb.org/	Trentino, Italy October 2014

Wider Public and Industry

The PARLANCE consortium members recognise the value and importance of disseminating project outcomes to the wider public and industry, both in terms of expected marketing benefits as well as recognition for PARLANCE partners as experts in the field and leaders in technology transfer.

Public Website

A public website is available at http://www.parlance-project.eu/. This was used as the main vehicle for dissemination and interaction with members of the public seeking information about the PARLANCE project. One important element for keeping the PARLANCE webpage interesting to visitors is that content is regularly updated and of high quality. Through the website, we provided public deliverables for download promptly after their finalisation, and news or other updates at least once per month. The website is also linked to our social media presence and has links to YouTube videos of demonstrations of the apps (with around 500 views).

Figures 3 and 4 show the statistics with regards the number of visits to the website by year and by quarter respectively. These figures reflect our increased dissemination efforts in the final year of the project almost doubling the number of website views as seen in Figure 3. In Figure 4, we can see that certain quarters have peaks of dissemination. In particular, the 3rd quarter of 2014 peaked with the SIGDIAL conference (technical co-chair and organiser Dr. Hastie) and the SemDial Conference (organised at Heriot-Watt by PARLANCE team members Dr. Verena Rieser and Dr. Arash Eshghi). Finally, Figure 5 shows the distribution of website users around the world. We can see here that the impact of the PARLANCE project is clearly international.

	Page Views	Unique Visits	First Time Visits	Returning Visits
Total	9,821	5,595	4,043	1,552
Average	2,456	1,399	1,011	388

	Page Views	Unique Visits	First Time Visits	Returning Visits
2014	4,348	2,520	1,985	535
2013	2,636	1,460	1,017	443
2012	2,676	1,489	938	551
2011	161	126	103	23

Figure 3: Visits to the PARLANCE project website, by year

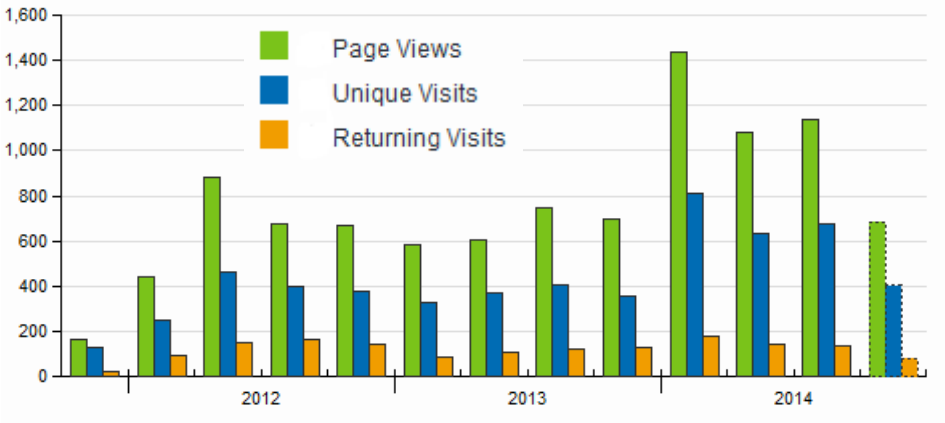

Figure 4: Visits to the PARLANCE project website, by quarter

Figure 5: PARLANCE project website visitor locations

Leaflet

A public leaflet describing the PARLANCE project has been designed in terms of content and layout (see Figure 6). This leaflet is used for the presentation of the project at major events such as conferences, workshops and fairs.

Figure 6: PARLANCE leaflet

Participation in Traditional and Social Media

Below are some of the highlights of participation in traditional media.

Article in *El Pais* Spanish newspaper

February 2012

Madrid, Spain

Radio programme on BBC Radio 4 – In Business Series

January 2013

Nationwide, UK


BBC NEWS TECHNOLOGY

Artificial intelligence: How to turn Siri into Samantha

By Leo Kelion
Technology writer

"Siri, why do you struggle with conversations?"

"I don't know what you mean - how about a web search for it?"

Related Stories

News article on the BBC website

February 2014

Nationwide, UK

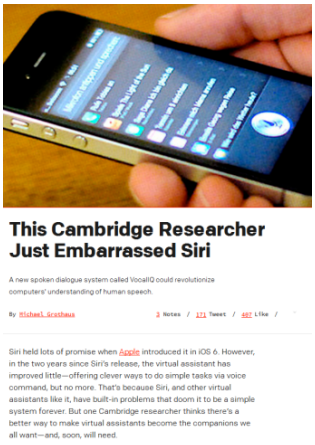
This Cambridge Researcher Just Embarrassed Siri

A new spoken dialogue system called VocalIQ could revolutionize computers' understanding of human speech.

Article on Fast Company (Co.Labs) website

February 2014

http://www.fastcolabs.com/3027067/this-cambridge-researcher-just-embarrassed-siri

Twitter, the social networking and microblogging site, is said to have more than 300 million users world-wide who follow, forward and respond to each other's 140-character tweets. In recent years, Twitter has become a new diffusion paradigm among researchers and innovators, reaching not only early adopters but also early majority adopters. For that reason, we set up a PARLANCE Twitter account, @parlanceproject, (http://twitter.com/parlanceproject) where we regularly tweeted about the advances in the project and related technology.

In addition to Twitter, we have set up a PARLANCE Facebook page (https://www.facebook.com/pages/Parlance-Project/277550145725220) and a LinkedIn Group (http://uk.linkedin.com/groups/PARLANCE-Project-5056353). These are public for all to join and are updated as per Twitter. Figure 7 shows the PARLANCE Twitter account homepage.

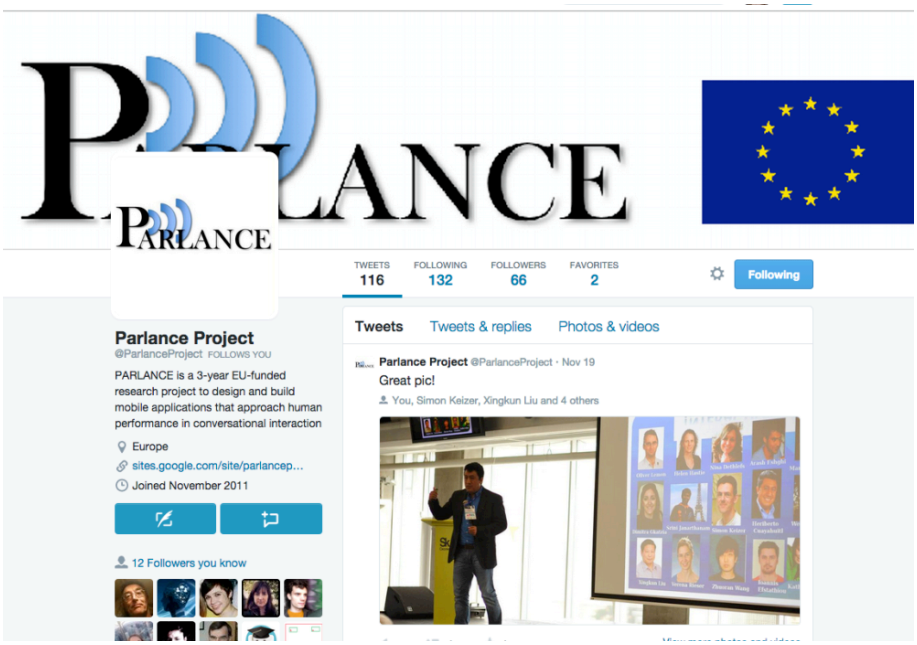

Figure 7: PARLANCE twitter account, November 2014

Parlance Mobile App

Android mobile applications were developed in the second and third year for dissemination in Spanish, English and Mandarin (See Figures 8 and 9). These were successfully demonstrated at a number of conferences including SIGDIAL 2013 (Metz, France) and 2014 (Philadelphia, USA) and SemDial 2014 (Edinburgh) with excellent feedback. At the end of the project, a more robust version of the English app was developed for the Google Play store under the title of 'Speak&Eat, San Francisco' (see Figure 10 for the Google Play store website).

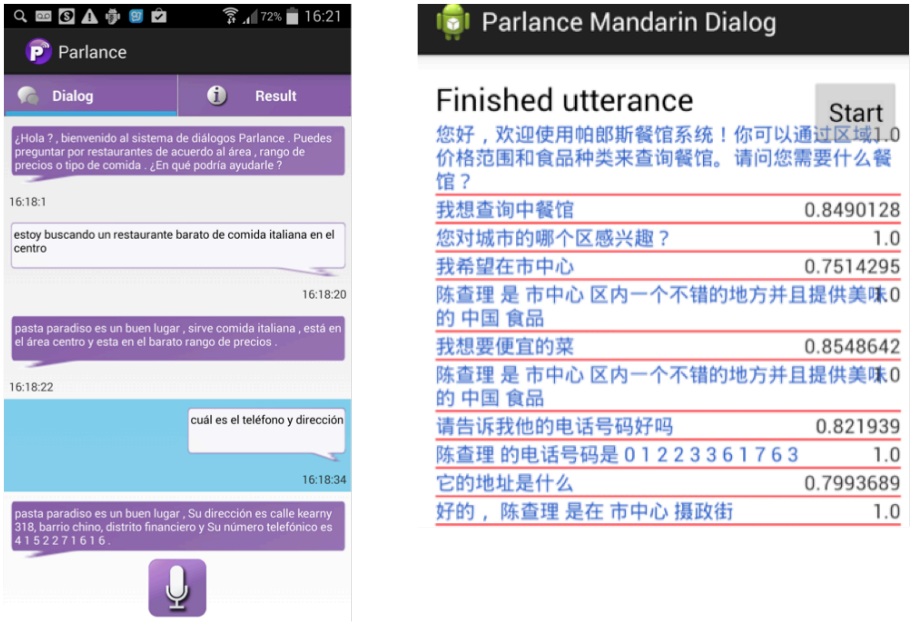

Figure 8: Spanish and Mandarin App for dissemination

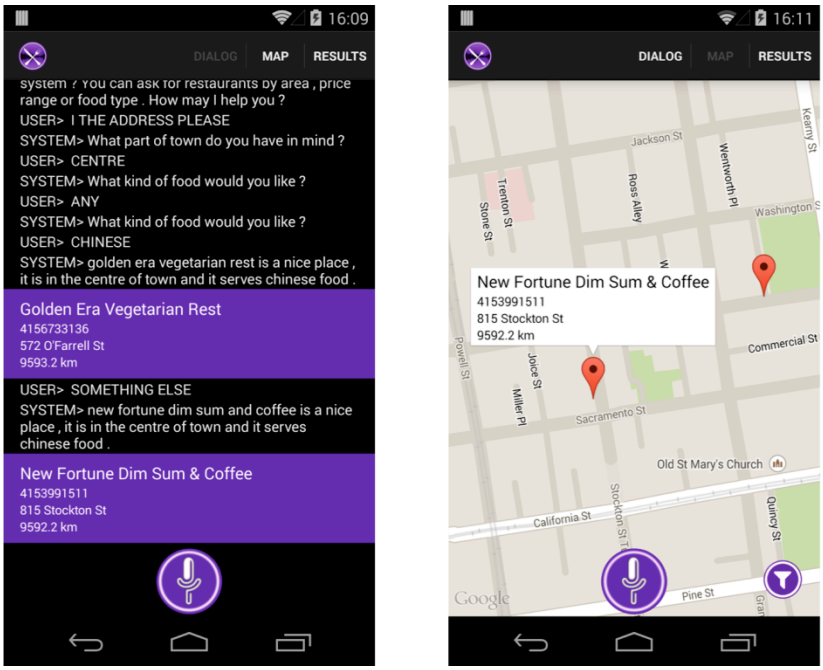

Figure 9: Final version of the English app

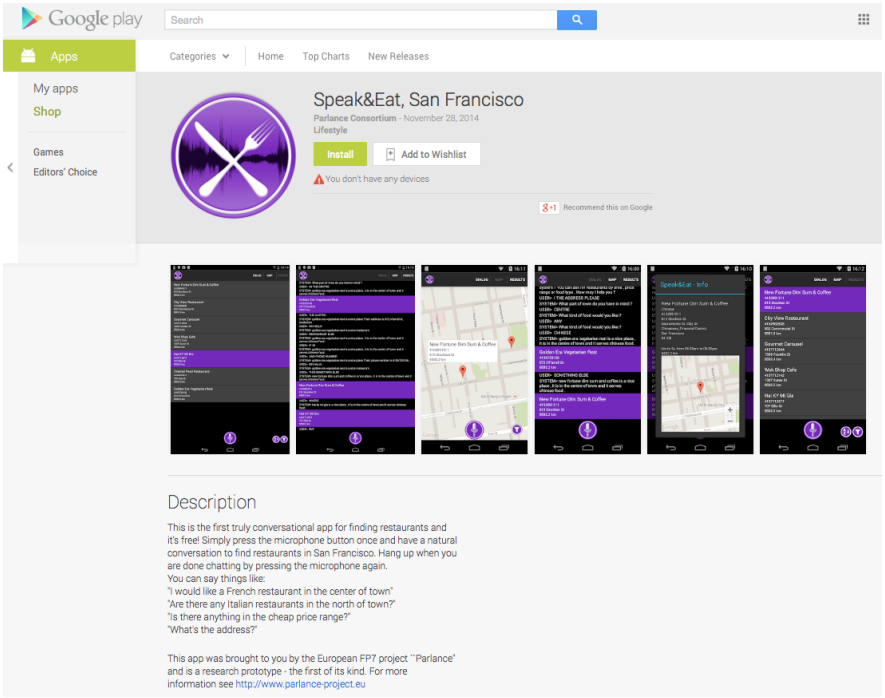

Figure 10: Screenshot of Speak&Eat, San Francisco App on the Google Play Store

Research Led Education

The academic partners have integrated and discussed the PARLANCE project in university courses encouraging further progress on the subject. This includes 'Interaction Design' course at Heriot-Watt where students designed and evaluated interactive products using speech technology inspired by PARLANCE including smart wheel chairs, interactive toys for sick children, interactive kitchen assistants etc. Dr Gašić has incorporated PARLANCE into a number of workshops and courses including:

- First ISCA summer school on speech processing in Crete for PhD students from across the Europe (August 2014), and
- Tutorial at Shanghai Jiao Tong University for final year undergraduates (July 2012).

Finally, Dr Hastie is on the advisory board of the REAL challenge, which inspires school pupils and undergraduates to design spoken dialogue systems. Dr Hastie attended the REAL challenge workshop in Baltimore in June 2014, where she spoke about PARLANCE.

Dissemination with Industry

As well as dissemination to the scientific community and the wider public, the consortium met and discussed PARLANCE with industry including BAE, Thales, Xerox, SeeByte, Microsoft, Google and others. These activities include 8 industry visits, 8 invited talks, 1 keynote talk and a number of presentations to industry. We hope that these activities will forge links and collaborations that will facilitate exploitation of the technology.

Site Leaders

Helen Hastie	h.hastie@hw.ac.uk
Steve Young	sjy@eng.cam.ac.uk
Peter Mika	pmika@yahoo-inc.com
Marie-Aude Aufaure	marie-aude.aufaure@ecp.fr
José Manuel Gómez Pérez	jmgomez@isoco.com
Jamie Henderson	james.henderson@unige.ch

Figure 11: Project Logo

Project website: http://www.PARLANCE-project.eu

Contact: Helen Hastie (Coordinator): h.hastie@hw.ac.uk

Section A (public)

LIST OF SCIENTIFIC (PEER REVIEWED) PUBLICATIONS, STARTING WITH THE MOST IMPORTANT ONES

No.	Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers[1]	Is/Will open access[2] provided to this publication?
1	POMDP-based dialogue manager adaptation to extended domains	MG, CB, MH, DK, MS, BT, PT &SY	Proceedings of SIGDIAL 2013	Annually	ACL	Metz, France	2013		Link on PARLANCE publications page	Yes through ACL Anthology
2	Context-Sensitive Natural Language Generation: From Knowledge-Driven to Data-Driven Techniques	ND, HC	Language and Linguistics Compass	Vol. 8, Issue 3, Mar 14	Wiley	New York, USA	2014	99-115	Link on PARLANCE publications page	No
3	Incremental on-line adaptation of POMDP-based dialogue managers to extended domains	MG, DK, PT, CB, MH, MS, BT, SY	Proceedings of INTERSPEECH 2014	Annually	ISCA	Singapore	2014		Link on PARLANCE publications page	Yes
4	Conditional Random Fields for Responsive Surface Realisation Using Global Features	ND, HH, HC & OL	Proceedings of the ACL 2013	Annually	ACL	Sofia, Bulgaria	2013		Link on PARLANCE publications page	Yes through ACL Anthology
5	Optimising Incremental Dialogue Decisions Using Information Density for Interactive Systems	ND, HH, VR, OL (HWU)	Proceedings of EMNLP-CoNLL 2012	Annually	ACL	Jeju Island, Korea	2012		Link on PARLANCE publications page	Yes through ACL Anthology

6	Word-Based Dialog State Tracking with Recurrent Neural Networks.	MH, BT, SY	Proceedings of SIGDIAL 2014	Annually	ACL	Philadelphia, USA	2014		Link on PARLANCE publications page	Yes through ACL Anthology
7	Hierarchical Reinforcement Learning for Situated Natural Language Generation	ND, HC	Natural Language Engineering	Jan-14	Cambridge University Press	Cambridge, UK	2014	n/a	Link on PARLANCE publications page	Yes
8	Policy Learning for Domain Selection in an Extensible Multi-domain Spoken Dialogue System.	ZW, HC, GW, HT, HW, HW	Proceedings of EMNLP 2014	Annually	ACL	Doha, Quatar	2014		Link on PARLANCE publications page	Yes through ACL Anthology
9	Ontology Learning from Open Linked Data and Web Snippets	IT, NM, YV, MAF (CRSA)	Proceedings of OTM 2012, OnToContent2012, Lecture Notes in Computer Science	Annually	Springer	Rome, Italy	2012	434—443	Link on PARLANCE publications page	Yes
10	Inverse Reinforcement Learning for Micro-Turn Management	DK, CB, PT, MG, MH, SY	Proceedings of INTERSPEECH 2014	Annually	ISCA	Singapore	2014		Link on PARLANCE publications page	Yes
11	Entity Recommendations in Web Search.	RB, BBC, PM, NT	Proceedings of ISWC 2013	Annually	On-line	Trentino, Italy	2013		Link on PARLANCE publications page	Yes
12	Dialogue Context Sensitive Speech Synthesis using Factorized Decision Trees	PT, CB, MG, MH, DK, SY	Proceedings of INTERSPEECH 2014	Annually	ISCA	Singapore	2014		Link on PARLANCE publications page	Yes

13	Metrics and Evaluation of Spoken Dialogue Systems	HH	Data-Driven Methods for Adaptive Spoken Dialogue Systems Computational Learning for Conversational Interfaces	Dec-12	Springer	New York	2012	131 - 150	Link on PARLANCE publications page	No
14	Continuous ASR for Flexible Incremental Dialogue	CB, MG, MH, DK, MS, BT, PT, KY, SY	Proceedings of ICASSP 2013	Annually	IEEE	Vancouver, Canada	2013		Link on PARLANCE publications page	Yes
15	Semantic Hyperlocal Search for PARLANCE Mobile Spoken Dialogue System	PA, MAA, NBM, HB, JMG, JH, BH, JL, PM & YV	Proceedings of ESWC 2013	Annually	Springer	New York, USA	2013		Link on PARLANCE publications page	Yes
16	Ontology-based User Preferences and Social Search for Spoken Dialogue Systems	YV, NM, MAF (CRSA)	Proceedings of 7th International Workshop on Semantic and Social Media Adaptation and Personalization	Annually	IEEE	Luxembourg	2012		Link on PARLANCE publications page	Yes
17	Dialogue Context Sensitive HMM-Based Speech Synthesis	PT, CB, MG, MH, DK, MS, BT, SY	Proceedings of ICASSP 2014	Annually	IEEE	Florence, Italy	2014		Link on PARLANCE publications page	Yes
18	Gaussian processes for POMDP-based dialogue manager optimization	MG, SY	IEEE Trans ASLP	Vol. 22, Issue 1	IEEE	N/A	2014		Link on PARLANCE publications page	No

19	Natural Language Generation as Incremental Planning Under Uncertainty: Adaptive Information Presentation for Statistical Dialogue Systems.	VR, OL, SK	IEEE/ACM Transactions on Audio, Speech and Language Processing	Vol. 22, Issue 5	IEEE	N/A	2014		Link on PARLANCE publications page	Yes
20	POMDP-based Statistical Spoken Dialogue Systems: a Review	SY, MG, BT, JW	Special Issue of the IEEE Proceedings on Speech Information Processing - Theory and Applications	Vol. 101, Issue 5	IEEE	N/A	2013	1160-1179	Link on PARLANCE publications page	Yes
21	Discriminative Spoken Language Understanding Using Word Confusion Networks	MH, MG, BT, PT, KY, SY	Proceedings of SLT 2012	Annually	IEEE	Florida,USA	2012		Link on PARLANCE publications page	Yes (with copyright notice)
22	The Use of Discriminative Belief Tracking in POMDP-based Dialogue Systems	DK, MH, MG, PT, SY	Proceedings of IEEE Spoken Language Technology Workshop (SLT), 2014.	Annually	IEEE	California, USA	2014		Link on PARLANCE publications page	Yes
23	Deep Neural Network Approach for the Dialog State Tracking Challenge	MH, BT, SY	Proceedings of SIGDIAL 2013	Annually	ACL	Philadelphia, PA, USA	2013		Link on PARLANCE publications page	Yes through ACL Anthology
24	Cluster-Based Prediction of User Ratings for Stylistic Surface Realisation	ND, HC, HH, VR, OL	Proceedings of EACL 2014	Annually	ACL	Gothenburg, Sweden	2014	702-711	Link on PARLANCE publications page	Yes through ACL Anthology

25	Training a Statistical Surface Realiser from Automatic Slot Labelling	HC, ND, HH, XL	To appear in Proceedings of SLT 2014	Annually	IEEE	California, USA	2014	N/A	Link on PARLANCE publications page	Yes (with copyright notice)
26	N-best Error Simulation, for training spoken dialogue systems	BT, MG, MH, PT, SY	Proceedings of SLT 2012	Annually	IEEE	Florida, U.S.A.	2012		Link on PARLANCE publications page	Yes (with copyright notice)
27	How domain-general can we be? Learning Incremental Dialogue Systems without Dialogue Acts	AE, OL	Proceedings of the Workshop on the Semantics and Pragmatics of Dialogue (SemDial)	Annually	On-line	Edinburgh, Scotland	2014		Link on PARLANCE publications page	Yes
28	Bridging the Gap Between Perception, Action and Communication	HC, LF, ND, MVO	Proceedings of the AAAI Workshop (MLIS)	Annually	AAAI Press	Quebec City, Canada	2014	N/A	Link on PARLANCE publications page	Yes
29	Towards Optimising Modality Allocation for Multimodal Output Generation in Incremental Dialogue	ND, VR, HH, OL	Proceedings of the AAAI Workshop (MLIS)	Annually	On-line	Montpellier, France	2012		Link on PARLANCE publications page	Yes
30	PQMPMS: A Preference-enabled Querying Mechanism for Personalized Mobile Search	BH, YV & MAA	Proceedings of RR 2013	Annually	On-line	Mannheim, Germany	2013		Link on PARLANCE publications page	Yes
31	Interactive Acquisition of Fuzzy Ontological Knowledge in Dialogue Systems	PA & JMG	Proceedings of K-CAP 2013	Annually	ACM	Banff, Canada	2013		Link on PARLANCE publications page	No

32	A Semi-Supervised Clustering Approach for Semantic Slot Labelling	HC, ND, HH	To appear in ICMLA 2014	Annually	IEEE	Detroit, MI, USA	2014	N/A	Link on PARLANCE publications page	Yes
33	Towards Ontology-Based Question Answering in Vague Domains	PA, MW	9th International Workshop on Semantic Media and Adaptation (SMAP) 2014	Annually	IEEE Computer Society	Corfu, Greece	2014		Link on PARLANCE publications page	To Appear
34	Dialogue Systems Using Online Learning: Beyond Empirical Methods	HC, ND	Proceedings of NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Community: Tools and Data	Annually	ACL	Montreal Canada	2012		Link on PARLANCE publications page	Yes through ACL Anthology
35	Collective Contexts in Conversation: Grounding by Proxy	AE, PGTH	To appear in Cognitive Science	Annually	Cognitive Science Society	N/A	2015		To appear	TBD
36	Graph-Based Seed Set Expansion for Relation Extraction Using Random Walk Hitting Times	JL & JH	Proceedings of NAACL HLT 2013	Annually	ACL	Georgia, USA	2013		Link on PARLANCE publications page	Yes through ACL Anthology
37	Policy optimisation of POMDP-based dialogue systems without state space compression	MG, MH, BT, PT, SY	Proceedings of SLT 2012	Annually	IEEE	Florida, U.S.A.	2012		Link on PARLANCE publications page	Yes (with copyright notice)
38	Sentiment Analysis for Dynamic User Preference Inference in Spoken Dialogue Systems.	YV, MC, ML, MA, ML	Proceedings of the First Workshop on Semantic Sentiment Analysis (SSA) at	Annually	On-line	Crete, Greece	2014		Link on PARLANCE publications page	Yes

			ESWC2014							
39	Barge-In effects in Bayesian Dialogue Act Recognition and Simulation	HC, ND, HH, OL	Proceedings of ASRU, 2013	Annually	On-line	Olumouc, Czech Republic	2013		Link on PARLANCE publications page	Yes
40	Optimizing Geographical Entity and Scope Resolution in Texts using Non-Geographical Semantic Information	PA, CR	Proceedings of the 6th International Conference on Advances in Semantic Processing (SEMAPRO 2012)	Annually	On-line	Barcelona, Spain	2012		Link on PARLANCE publications page	Yes
41	Towards Learning Personalized Semantic Relevance Paths In Dialogue Systems	PA, MW	Proceedings of the 8th International Workshop on Semantic Media and Adaptation	Annually	On-line	Bayonne, France	2013		Link on PARLANCE publications page	No
42	Optimising Incremental Generation for Spoken Dialogue Systems: Reducing the Need for Fillers	ND, HH, VR, OL	Proceedings of INLG 2012	Annually	ACL	Starved Rock, IL, USA	2012		Link on PARLANCE publications page	Yes through ACL Anthology
43	Context-based Grouping and Recommendation in MANETs	YV, MK, NBM & MAA	Intelligent Technologies and Techniques for Pervasive Computing, IGI Global	N/A	IGI Global	USA	2013	157 - 178	Link on PARLANCE publications page	No

44	Scenario-Driven Selection and Exploitation of Semantic Data for Optimal Named Entity Disambiguation.	PA, CR, JMG	Proceedings of the Semantic Web and Information Extraction Workshop (SWAIE 2012)	Annually	On-line	Galway, Ireland	2012		Link on PARLANCE publications page	Yes
45	Machine Learning for Interactive Systems and Robots: A Brief Introduction	HC, MO, ND & LF	Proceedings of MLIS 2013	Annually	ACM	Beijing, China	2013		Link on PARLANCE publications page	Yes
46	Getting to Know Users: Accounting for the Variability in User Ratings.	ND, HC, HH, VR, OL	Abstract paper in the Workshop on the Semantics and Pragmatics of Dialogue (SemDial) 2014	Annually	On-line	Edinburgh, Scotland	2014	107-172	Link on PARLANCE publications page	Yes
47	Impact of ASR N-Best Information on Bayesian Dialogue Act Recognition	HC, ND, HH & OL	Proceedings of SIGDIAL 2013	Annually	ACL	Metz, France	2013		Link on PARLANCE publications page	Yes through ACL Anthology
48	Two Alternative Frameworks for Deploying Spoken Dialogue Systems to Mobile Platforms for Evaluation “in the Wild”.	HH, MA, PA, HB, HC, ND, MG, JH, OL, XL, PM, TP, VR, PT, YV, BVT, MY, SY, YY	Abstract paper in the Workshop on the Semantics and Pragmatics of Dialogue (SemDial), 2014	Annually	On-line	Edinburgh, UK	2014		Link on PARLANCE publications page	Yes
49	Mining the Web for Points of Interest	AR, VM, AP, HB (YAHOO!)	Proceedings of SIGIR, 2012	Annually	ACM	Portland, Oregon	2012		Link on PARLANCE publications page	Yes

50	Comparing HMMs and Bayesian Networks for Surface Realisation	ND, HC	Proceedings of NAACL-HLT, 2012	Annually	ACL	Montreal, Canada	2012		Link on PARLANCE publications page	Yes through ACL Anthology
51	The PARLANCE mobile app for interactive search in English and Mandarin.	HH, MA, PA, HB, HC, ND, MG, JH, OL, XL, PM, NBM, TP, VR, BT, PT, VY, BV, MY, SY, YY	Proceedings of SIGDIAL 2014	Annually	ACL	Philadelphia, USA,	2014		Link on PARLANCE publications page	Yes through ACL Anthology
52	Incremental spoken dialogue systems: Tools and data.	HH, OL, ND	Proceedings of NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Community: Tools and Data	Annually	ACL	Montreal Canada	2012		Link on PARLANCE publications page	Yes through ACL Anthology
53	Towards Comparative Evaluation and Shared Tasks for NLG in Interactive Systems.	AB, HH	In Srinivas Bangalore and Amanda Stent (eds.), Natural Language Generation in Interactive Systems.	N/A	Cambridge University Press	N/A	2013		Link on PARLANCE publications page	No
54	On-line Policy Optimisation of Bayesian Spoken Dialogue Systems via Human	MG, CB, MH, DK, MS, BT, PT &SY	Proceedings of ICASSP 2013	Annually	IEEE	Vancouver, Canada	2013		Link on PARLANCE publications page	Yes

	Interaction.									
55	Efficient Computation of Mean Truncated Hitting Times on Very Large Graphs	JL, JH	In Journal CoRR, Cornell University				2013		Link on PARLANCE publications page	Yes
56	Demonstration of the PARLANCE system: a data-driven incremental, spoken dialogue system for interactive search	HH, MAA, PA, HC, ND, MG, JH, OL, XL, PM, NBM, VR, BT, PT & YV	Proceedings of SIGDIAL 2013	Annually	ACL	Metz, France	2013		Link on PARLANCE publications page	Yes through ACL Anthology
57	Optimising Incremental Generation for Information Presentation of Mobile Search Results	ND, HH, VR, OL	Symposium: Influencing people with information (SIPI)	One off	On-line	Aberdeen, UK	2012		Link on PARLANCE publications page	Yes through SIPI

LIST OF DISSEMINATION ACTIVITIES (IN CHRONOLOGICAL ORDER)							
No	**Type of activities**	**Title**	**Date**	**Place**	**Type of audience**	**Size of audience**	**Countries addressed**
1	Website	The website is updated regularly to include any recent, relevant information including publications.	Nov 11 to Date	N/A	General Public	Unknown	International
2	Social Media	Set up of Facebook and LinkedIn pages, linked to Twitter and updated regularly with recent, relevant news.	Nov 11 to Date	N/A	General Public	Unknown	International
3	Invited Talk	Presentation at SICSA workshop on mobile interaction (Nina Dethlefs)	Jan-12	Glasgow, UK	Scientific community	50	UK
4	Press release	Article in the *El País* Spanish newspaper	Feb-12	Madrid, Spain	General public	350,000	Spain
5	Invited Talk	Presentation at SICSA workshop on multimodal interaction (Oliver Lemon)	Feb-12	Abertay, UK	Scientific community	50	UK
6	Invited Talk	"User Modelling and modular ontological knowledge base for mobile conversational search," Nesrine Ben Mustapha and Yves Vanrompay, EURECOM-CEA-MAS seminar	Mar-12	Paris, France	Industry and scientific community	22	International
7	Web	Parlance on LT-Innovate website (http://lt-innovate.eu/page/eu-projects);	Mar-12	N/A	Policy makers, industry and scientific community		Europe
8	Press release	Article in University of Cambridge on-line magazine, Research Features (http://www.cam.ac.uk/research/features/call-of-the-wired/)	Apr-12	Cambridge, UK	General Public		UK
9	Invited Talk	Adam Rae "Uncovering Region Boundaries within Geographic Data." and Hugues Bouchard "Hyper-local Trend Detection using Social Media." Where-Camp EU	Apr-12	Amsterdam, The Netherlands	Scientific community		Europe

10	Invited Talk	"Computational Geography." Invited Talk at Bing, Bellevue WA. (Vanessa Murdock)	May-12	Bellevue WA	Industry and scientific community		USA
11	Invited Talk	"Modelling Locations with Social Media." At RMIT University (Vanessa Murdock)	May-12	Melbourne, Australia	Scientific community		Australia
12	Web	The consortium joined META (http://www.meta-net.eu);	May-12	N/A	Policy makers, industry and scientific community		Europe
13	Press release	Article in the Cambridge News	Jun-12	Cambridge, UK	General public	100, 000	UK
14	Invited Talk	Knowledge harvesting for Business Intelligence: Nesrine Ben Mustapha and Marie-Aude Aufaure, Summer School EBISS	Jul-12	Brussels, Belgium	Industry and scientific community	42	International
15	Workshop Organisation	Co-organized the 1st Joint International Workshop on Entity-oriented and Semantic Search (JIWES) at SIGIR 2012	Aug-12	Portland, Oregon, USA	Scientific community	50	International
16	Workshop Organisation	Organised MLIS, Machine Learning for Interactive Systems at ECAI (Heriberto Cuayahuitl)	Aug-12	Montpellier, France	Scientific community	16	International
17	Keynote Talk	"Recent Developments in Statistical Spoken Dialogue Systems." By Prof. Steve Young at "Eurosip Seminar at the 10th ITG Conference on Speech Communication	Sep-12	Braunschweig , Germany	Industry and scientific community	130	International
18	Workshop Organisation	Co-organised and sponsoring the 5th Intl. Workshop on Exploiting Semantic Annotations in Information Retrieval (ESAIR) at CIKM 2012 (Peter Mika)	Oct-12	Hawaii, USA	Scientific community	30	International
19	Keynote Talk	Keynote at the 3rd Russian Semantic Web Conference (Peter Mika)	Oct-12	St Petersburg, Russia	Scientific community	80	International
20	Invited Talk	Invited talk at Yandex, St Petersburg. Making the web searchable (Peter Mika)	Oct-12	St Petersburg, Russia	Industry and scientific	120	Russia

					community		
21	Invited Talk	Presentation at Architectures for Conversationally Competent Spoken Dialogue Systems (ACCSDS) workshop (Nina Dethlefs)	Oct-12	Hamburg, Germany	Scientific community	40	International
22	Flyer	Parlance Flyer	Nov-12	N/A	Industry and scientific community		International
23	Fact sheet and overview slides	Fact sheet and overview slides	Nov-12	N/A	Industry and scientific community		International
24	Invited Talk	'Statistical Dialogue Optimisation, Inference and Control for Complex Dynamical Systems workshop' by Milica Gasic at the Centre for Complexity, University of Warwick	Nov-12	Warwick, USA	Scientific Community	30	UK
25	Radio Programme	'In Business' by Steve Young on BBC Radio 4	Jan-13	UK	General Public	Unknown	International
26	Invited Talk	'Reinforcement Learning of Adaptive Dialogue Systems' by Verena Rieser at Electrical Engineering Dept., Technical University Munich	Feb-13	Munich, Germany	Scientific Community	20	Germany
27	Invited Talk	'Planning Under Uncertainty for Output Generation' by Verena Rieser at Computer Lab, Cambridge University	Mar-13	Cambridge, UK	Scientific Community	20	UK
28	Invited Talk	'On-line Dialogue Policy Optimisation' by Milica Gasic at the Computer Science & Artificial Intelligence Lab (CSAIL), MIT	May-13	MA, USA	Scientific Community	30	USA
29	Invited Talk	'Statistical Learning for Natural Conversational Interfaces' by Oliver Lemon at Nuance Communications Inc.	May-13	California, USA	Industry	30	USA
30	Keynote Talk	'Related entity finding on the Web' by Peter Mika at WWW 2013 workshop WOLD	May-13	Rio de Janerio, Brazil	Scientific community	60	International

31	Guest lecture	'Related entity finding on the We' by Peter Mika at IT4BI, University of Tours	May-13	Tours, France	Scientific community	20	International
32	Invited Talk	'User Modelling and personalization for spoken search systems' by Yves Vanrompay at Ecole Centrale, Paris	Jun-13	Paris, France	Scientific community	14	French
33	Tutorial	'Introduction to Semantic Search' by Peter Mika at the Yahoo! Summer School on Semantic Search	Jun-13	Bangalore, India	Scientific community	250	India
34	Invited Talk	'Optimising Conversational Interfaces' by Verena Rieser at ICSI, Berkley	Jul-13	Berkley, USA	Scientific Community	12	Finland, Germany, Singapore
35	Invited Talk	'Related entity finding on the Web' by Peter Mika at SSSW'13	Jul-13	Madrid, Spain	Scientific community	50	International
36	Invited Talk	'Semantic Search' by Peter Mika at Stanford BMIR	Aug-13	California, USA	Scientific community	20	USA
37	Workshop Organisation	Bridging the Gap Between Perception, Action and Communication at IJCAI Workshop on Machine Learning for Interactive Systems - MLIS 2013 (Heriberto Cuayahuitl)	Aug-13	Beijing, China	Scientific Community	32	International
38	Industry visit	British Gas Visit by Helen Hastie	Aug-13	Edinburgh, UK	Industry	3	UK, South America
39	Keynote Talk	'Web-scale Spoken Dialogue Systems' by Steve Young at EUSIPCO 2013	Sep-13	Marrakech, Morocco	Scientific Community	300	International
40	Keynote Talk	'Exploiting Information Extraction and the Semantic Web at Yahoo! Search' by Peter Mika at ISWC 2013 LD4IE workshop	Oct-13	Sydney, Australia	Scientific community	40	International
41	Inaugural Lecture	'Towards Open-domain Spoken Dialogue Systems' by Steve Young at CSAIL Human Language Technology Distinguished Lecture Series	Oct-13	MA, USA	Scientific Community	100	International
42	Keynote Talk	What's Up in Business Intelligence? A Contextual and Knowledge-based Perspective (ER conference 2013) 32nd International Conference on Conceptual Modelling (Marie-	Nov-13	Hong Kong	Scientific Community	100	International

		Aude Aufaure)					
43	Invited Talk	Multi-Task Learning for Interactive Systems and Robots - LORIA, INRIA (Heriberto Cuayahuitl)	Nov-13	Nancy, France	Scientific Community	30	France
44	Invited Talk	Statistical Dialogue Optimisation using Gaussian Processes, conference "35 years of Computational Linguistics in Serbia" - University of Belgrade (Milica Gasic)	Nov-13	Belgrade, Serbia	Scientific Community	30	Serbia, Bulgaria
45	Inaugural Lecture	Inaugural lecture on Human Language Technology - MIT/Apple (Steve Young) https://itunes.apple.com/us/itunes-u/human-language-technology	Dec-13	Online	Scientific Community	Unknown	USA
46	Invited Talk	How to write and EC proposal - Scottish Parliament (Helen Hastie)	Jan-14	Edinburgh, UK	Scientific Community	40	UK
47	Keynote Talk	Hierarchical Reinforcement Learning for Interactive Systems and Robots - AINL 2014 (Heriberto Cuayahuitl)	Jan-14	Moscow, Russia	Scientific Community	150	Russia
48	TV Interview	Virtual Personal Assistants on BBC Newsnight (Steve Young)	Feb-14	UK	General Public	Unknown	UK
49	News Article	Artificial Intelligence: How to turn Siri into Samantha http://www.bbc.co.uk/news/technology-26147990	Feb-14	UK	General Public	Unknown	UK
50	News Article	Fast Company (Co.Labs) article http://www.fastcolabs.com/3027067/this-cambridge-researcher-just-embarrassed-siri	Feb-14	UK	General Public	Unknown	UK
51	Invited Talk	PARLANCE. NLP for User Generated Content workshop, Yahoo Labs Winter Science Week (Peter Mika)	Mar-14	Sunnyvale, CA, USA	Industry	30	USA
52	Guest lecture	Entity Recommendations in Web Search lecture - Centruum voor Wiskunde en Informatica –CWI (Peter Mika)	Apr-14	Amsterdam, Netherlands	Scientific Community	20	Netherlands
53	Invited Talk	Semantic Search at Yahoo. The 36th European Conference on Information Retrieval (ECIR 2014) Industry Day. http://ecir2014.org/industry-day/ (Peter Mika)	Apr-14	Amsterdam, Netherlands	Scientific Community	80	Europe

54	Industry Visit	Thales visit by Helen Hastie	May-14	London, UK	Industry	4	UK
55	Conference Organisation	SIGDIAL Conference 2014 (Helen Hastie – Technical Chair)	Jun-14	Philadelphia, USA	Scientific Community	120	International
56	Keynote Talk	Statistical Approaches to Open-domain Spoken Dialogue Systems (Steve Young)	Jun-14	Philadelphia, USA	Scientific Community	120	International
57	Workshop Organisation	REAL challenge workshop (Helen Hastie)	Jun-14	Baltimore, USA	Scientific Community	40	International
58	Keynote Talk	Semantic Search. Keynote at the 3rd Spanish Conference on Information Retrieval (CERI 2014). http://ceri2014.udc.es/invited_speakers (Peter Mika)	Jun-14	A Coruña, Spain	Scientific Community	30	Spain
59	Invited Talk	Presentation to the Ministry of Defence at the ASUR conference (Helen Hastie)	Jul-14	Winchester, UK	Industry	40	UK
60	Workshop Organisation	Workshop on Machine Learning for Interactive Systems (MLIS) 2014 (Heriberto Cuayahuitl)	Jul-14	Quebec, Canada	Scientific Community	40	Canada
61	Industry Visit	Presentation and meeting - British Gas, Rio group (Helen Hastie)	Aug-14	Edinburgh, UK	Industry	1	Brazil
62	Invited Talk	Overview of Research Projects within iSOCO (Boris Villazón-Terrazas)	Aug-14	Cuenca, Ecuador	Scientific Community	30	Europe
63	Invited Talk	Overview of Research Projects within iSOCO (Boris Villazón-Terrazas)	Aug-14	Loja, Ecuador	Scientific Community	20	Europe
64	Invited Talk	Statistical Spoken Dialogue Systems - ISCA Summer school on Speech Processing, University of Crete (Milica Gasic)	Aug-14	Heraklion, Greece	Scientific community	30	Europe
65	Invited Talk	Semantic Search on the Rise. 10th Semantic Technology and Business Conference. http://semtechbizsj2014.semanticweb.com/agenda.cfm?confid=82&scheduleDay=08/20/14 (Peter Mika)	Aug-14	San Jose, CA, USA	Industry	50	USA

66	Invited Talk	Overview of Interaction Lab - Centre for Robotics (Oliver Lemon)	Sep-14	Edinburgh, UK	Scientific Community	70	UK
67	Keynote Talk	Challenges and Opportunities Induced by Big Data and Open Data for Business Intelligence at IEEE CIST 2014 (Marie-Aude Aufaure)	Oct-14	Tetouan-Cherfchaouen, Morocco	Scientific Community	100	International
68	Workshop Organisation	Second year students partake in REAL challenge (Helen Hastie)	Oct-14	Edinburgh, UK	Scientific Community	80	UK
69	Invited Talk	Statistical Spoken Dialogue Systems and their adaptation to changes in the dialogue domain - Microsoft Research Cambridge (Steve Young)	Oct-14	Cambridge, UK	Scientific Community	15	UK