Towards teachable dialogue systems

We have developed an approach to mixed-initiative, flexible dialogues for interface agents[1] in which dialogue moves are generated by the agent under the influence of a *task model*[2] and an *extensible concept ontology*[3]. The task specification and knowledge base provide a general context for the user-agent dialogue, including guidance and constraints in responding appropriately to unscripted user input. The components of this system are based on models of human decision-making and planning, ontological models of medical expertise and current dialogue theory. The work is therefore relevant to the proposed grand challenge "Computational modelling of human language function".

The system is currently undergoing technical evaluation and usability testing. Technical evaluations include interpretive accuracy and scalability (in terms of vocabulary and task repertoire) and usability tests include naturalness and flexibility of the dialogue. This work is being carried out in the context of the CREDO project, a major development at Cancer Research UK aimed at introducing computerised decision support and care plan management in the treatment of women with breast cancer. A scenario in which a GP telephones for advice on whether a patient requires urgent referral for suspected breast cancer is demonstrated at http://www.openclinical.org/dm_homey.html.[4]

In most practical dialogue systems, the dialogue designer specifies the exact interactions that can take place in an application. This hand coding allows precise control of what can occur within a dialogue but it is an expensive process, especially for complex dialogues where the number of states can be large. In the medical domain, where knowledge bases are particularly complex, the number of states could run into millions. Moreover, this kind of approach leads to inflexible dialogues, and is not appropriate if the task is changing dynamically (for example, during the enactment of a care plan that may be active for weeks or months and patient circumstances and clinical needs are changing). We therefore wish to initiate a further phase of research in which the system can acquire new plans and concepts by spoken instruction.

In our view this may now be within reach. For example the CREDO the system may know that a particular clinical goal needs to be achieved at a particular point in a plan, but not how to achieve it. In such situations a clinician could instruct the system in the tasks to be carried out (e.g. "give 50 mls of CMF, record the white blood count and book a mammogram") and the agent would generate the appropriate task specification for later enactment. During such instruction it may be necessary for the system to learn new concepts, their relationships with concepts already present in the domain ontology, and new terms for concepts. These tasks are currently carried out by hand but implementation via spoken instruction should be tractable.

The long term goal of this project will be to demonstrate higher levels of dialogue function and autonomy of interface agents than is currently possible. This would have theoretical benefits in increasing our understanding of conversational strategies, and practical benefits in that interface agents can expand their conversational repertoire by being taught dialogue plans during use.

Martin Beveridge and John Fox, Cancer Research UK

---

[1] EU 5th Framework Project HOMEY

[2] Using the PRO*forma* task specification language and agent design tools (Fox and Das, *Safe and Sound: Artificial Intelligence in Hazardous Applications*, AAAI and MIT Press, 2000)

[3] Developed with the assistance of Language and Computing nv, Belgium.

[4] The knowledge and medical vocabulary are somewhat specialised but the domain independent capabilities of the dialogue manager should be evident in this demonstration.