# Current research topics in spoken dialogue systems

Milica Gašić

Dialogue Systems Group, Cambridge University Engineering Department

March 3, 2016

Open domain dialogue systems

Incremental dialogue

# Statistical spoken dialogue systems

- So far we focused on limited-domain spoken dialogue systems
- These are slot filling dialogues of relatively small ontologies $\sim$ 1000 entities $\sim$ 20 slots
- Provided we have data each dialogue system module can be described with a statistical model and trained

# Dialogue system architecture

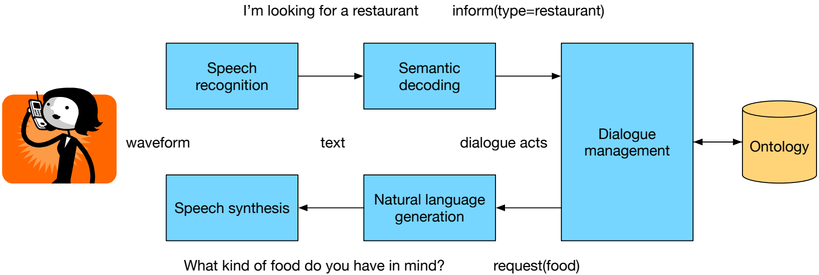

# Ontology

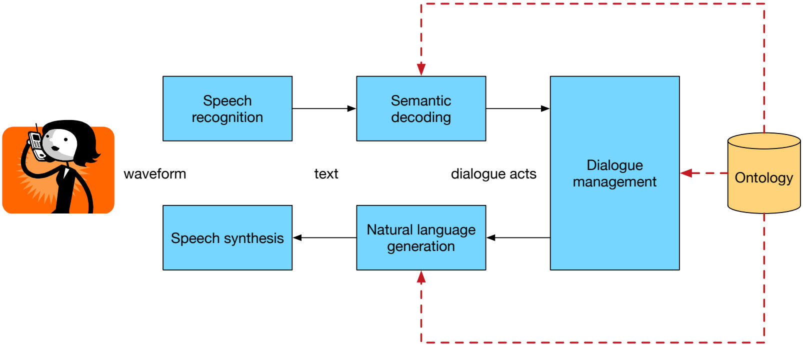

Speech recognition
Semantic decoding
Dialogue management
Ontology
waveform
text
dialogue acts
Speech synthesis
Natural language generation

# Extending dialogue system coverage

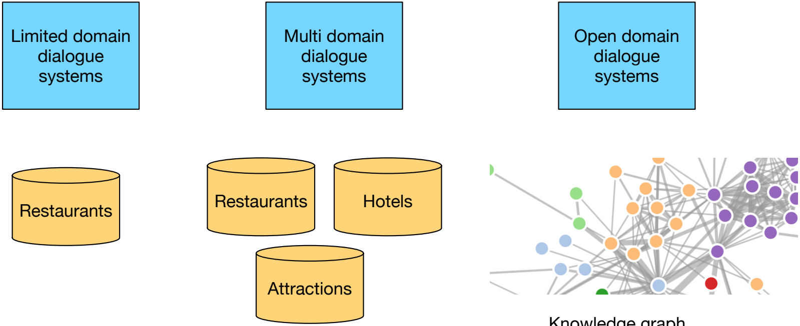


Knowledge graph

# Objectives

- Allow the user to reference multiple domains within a single conversation
- Support natural conversation even for domains which have been rarely visited
- Learns automatically through interaction with the user

Deploy

Collect data

Improve

# Example dialogue

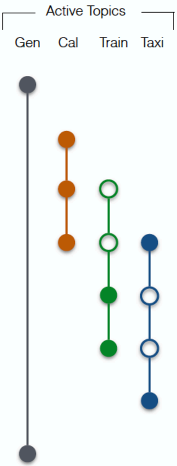


**Hello, how can I help you?**

*What appointments do I have tomorrow?*

**You have a meeting at 10am with John and a teleconf at noon with Bill.**

*I need to go to London first thing, can you reschedule the meeting with John?*

**John is free tomorrow at 3pm, is that ok?**

*Yes, thats fine. I also need a taxi to the station.*

**Meeting with John at 15.00 is confirmed. What time do you need the taxi?**

*When does the train depart to London?*

**The 9.15am gets in at 10.06.**

*When is the one before that?*

**The train before that leaves at 8.45am and arrives at 9.40.**

*Ok I will take that, book the taxi for 8.15am from my house.*

**Ok, I will book the taxi for 8.15am, is that correct?**

*Yes that's right.*

**Ok. Do you need anything else?**

*Not for now thanks.*

# Reuse of data

- Each element of a dialogue system is data driven and depends on the ontology
- If we want to build a dialogue system to operate on multiple domains it is essential there is reuse of knowledge

# Delexicalisation for domain-independent belief tracking

- Idea: replace all occurrences of slot values in data with generic symbols (tags) and train a general belief tracker

*I want Chinese food*
*I want cheap price range* ⟹ *I want* VALUE SLOT

- This captures dialogue dynamics and facilitate generalisation to unseen slots and domains

# Initialising belief tracking models for new domains [Mrkšić et al., 2015]

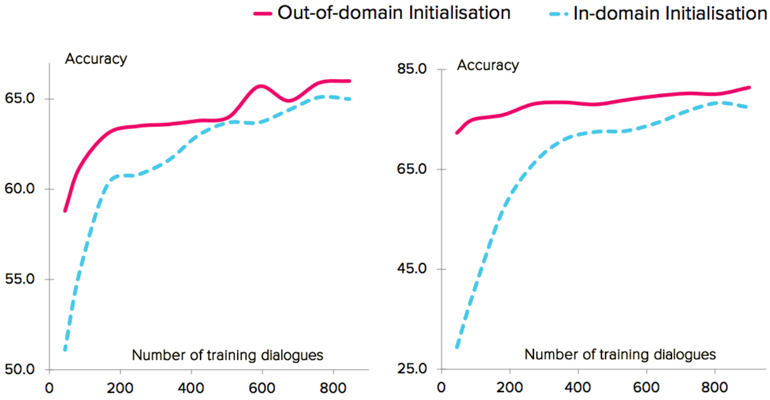


In-domain: Michigan Restaurants
Out-of-domain: Other restaurant domains, Tourist Information, Hotels

In-domain: Laptops
Out-of-domain: Restaurants, Tourism, Hotels

# Committee of dialogue policies

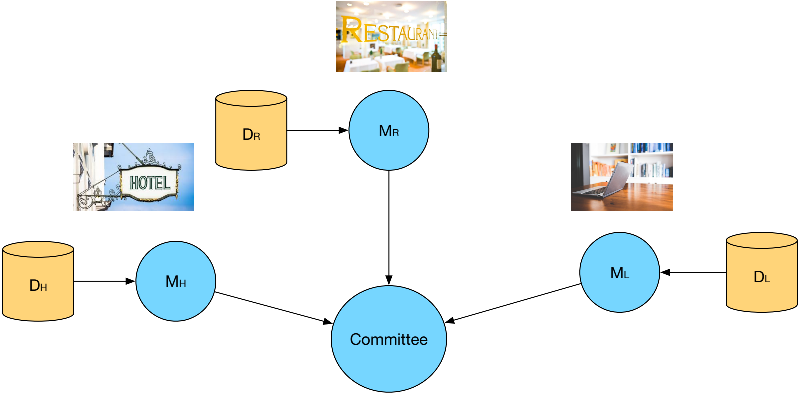

# Bayesian committee machine for distributed dialogue management

- Each dialogue manager operates independently, receives speech, tracks its own beliefs and proposes system actions
- Each domain policy is modelled as a Bayesian estimator, e.g. Gaussian process with posterior mean $\overline{Q}_i$ and covariance $\Sigma_i$
- Committee of dialogue policies is then also a Bayesian estimator

$$\overline{Q}(\boldsymbol{b}, a) = \Sigma(\boldsymbol{b}, a) \sum_{i=1}^{M} \Sigma_i(\boldsymbol{b}, a)^{-1} \overline{Q}_i(\boldsymbol{b}, a)$$

$$\Sigma(\boldsymbol{b}, a)^{-1} = -(M-1) * k((\boldsymbol{b}, a), (\boldsymbol{b}, a))^{-1} + \sum_{i=1}^{M} \Sigma_i(\boldsymbol{b}, a)^{-1}$$

# Results [Gašić et al., 2015]

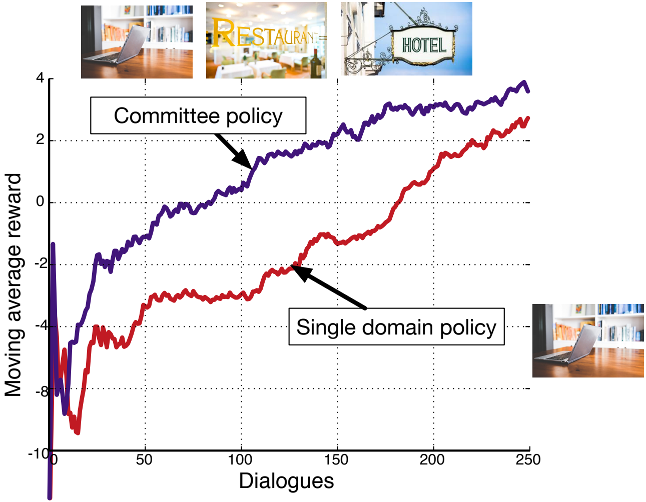

# But that is not all...

- Semantic decoder needs to share data for new domains
- Dialogue manager must be able to talk about different things at the same time
- Natural language generator needs also to be able to generate responses for new domains
- ...

# Towards more natural interaction

- So far we assumed strict turn taking mechanism
- Humans interrupt each other, talk at the same time and hesitate during conversation

U: I'm looking for an Italian restaurant in the north

S: Cocum is a nice Indian restaurant in the north

U: Italian, please

S: Margarita is a nice Italian restaurant

U: I'm looking for umm…

S: Yes?

The user is hesitates — the system issues a **backchannel**

U: …an Italian restaurant in the north

S: Cocum is a nice Indian …

U: Italian, please

The user interrupts to correct misunderstanding — **barge-in**

S: Margarita is a nice Italian restaurant

# Incremental dialogue

Incremental dialogue architecture allow next component to start processing before the previous one has completed.

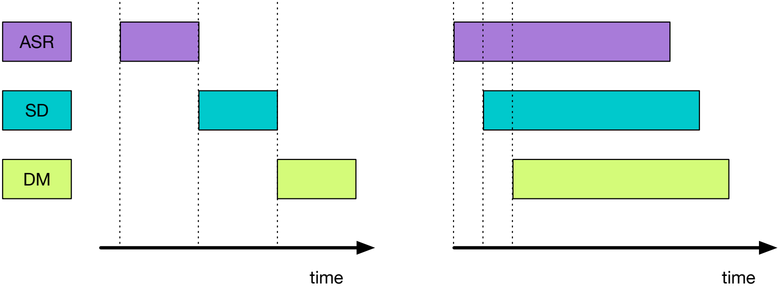

# Why incremental dialogue?

The main benefits of incremental dialogue are:

- reduced latency
- increased system efficiency
- increased user satisfaction
- increase user fluency

Open questions

- What is the best architecture?
- Which machine learning methods to use to find optimal behaviour?

# Incremental speech recognition and semantic decoder [Paetzel et al., 2015]

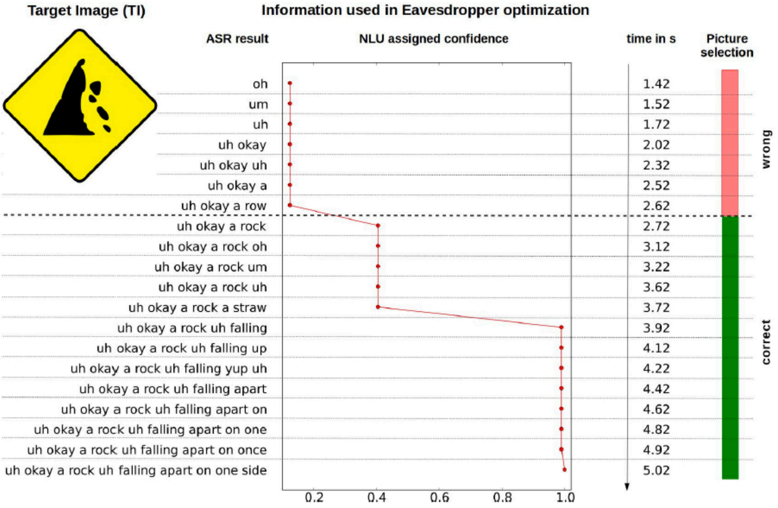

# Incremental dialogue manager

Incremental dialogue manager apart from making high level decisions also needs to make low level decisions (micro-turns):

- speak
- be silent
- back-channel

How to learn these decisions?

# Learning micro-turn decisions

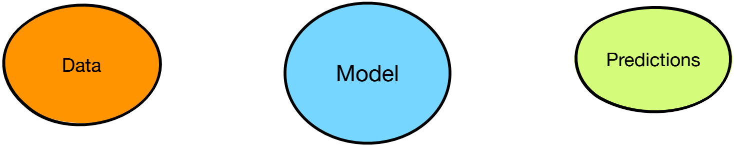


- States and decisions
- States and rewards
- States and decisions

- Supervised learning model
- Reinforcement learning model
- Inverse reinforcement learning model

- Decisions that mimic the data
- Decisions that maximise the reward
- Decisions that mimic the data and maximise the reward

# Supervised vs inverse reinforcement learning

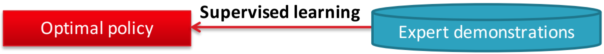


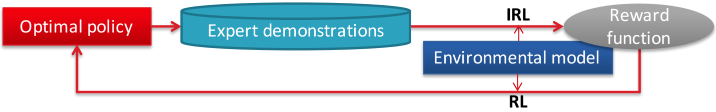


- Inverse reinforcement learning (IRL) recovers the model and then learns the optimal policy

# Micro-turn management as an IRL task [Kim et al., 2014]

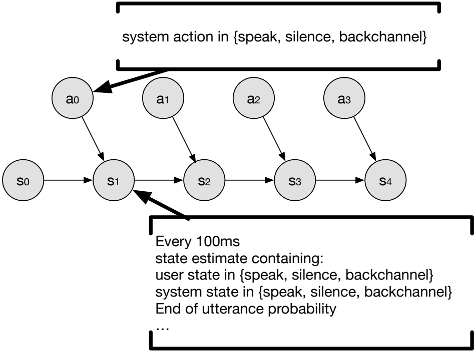


- Learn optimal decisions from human-human data by using Gaussian process IRL which assumes a non-parametric (and non-linear) model for the unknown reward.

# Difficulty of evaluation

What to measure?

- Log likelihood of the micro-turn
- Length of dialogue
- User satisfaction

Remaining problems

- Effectively incorporate semantics into the incremental decision making process

# IRL for micro-turn management – results

- Cross-validation log-likelihood results of each micro-turn
- Reward can depend only on state $R(s)$ or both state and action $R(s, a)$
- Supervised learning for $\gamma = 0$

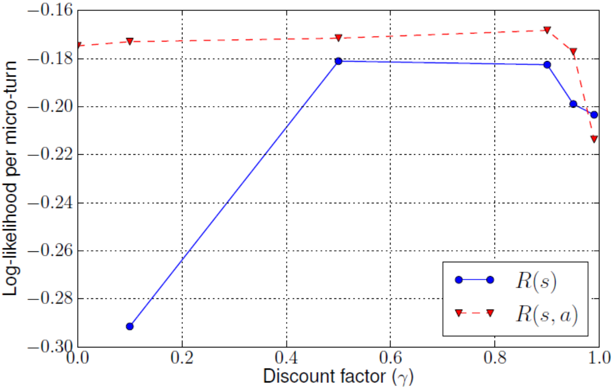

# Summary

- End-to-end statistical dialogue is feasible for and can match or exceed hand-crafted systems in limited domains
- Reuse of data is crucial for modelling open-domain systems and avoiding unacceptable user experience
- Incremental dialogue produces more natural interaction by increasing system efficiency, user satisfaction and user fluency

# References I


Gašić, M., Mrkšić, N., Su, P., Vandyke, D., Wen, T., and Young, S. (2015).
Policy committee for adaptation in multi-domain spoken dialogue systems.

Kim, D., Breslin, C., Tsiakoulis, P., Gasic, M., Henderson, M., and Young, S. J. (2014).
Inverse reinforcement learning for micro-turn management.
In *INTERSPEECH*, pages 328–332.

Mrkšić, N., Séaghdha, D. Ó., Thomson, B., Gašić, M., Su, P.-H., Vandyke, D., Wen, T.-H., and Young, S. (2015).
Multi-domain dialog state tracking using recurrent neural networks.
*arXiv preprint arXiv:1506.07190*.

# References II

Paetzel, M., Manuvinakurike, R., and DeVault, D. (2015).
"so, which one is it?" the effect of alternative incremental architectures in a high-performance game-playing agent.
In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 77–86, Prague, Czech Republic. Association for Computational Linguistics.