Dialogue management: generative approaches to belief tracking

Milica Gašić

Dialogue Systems Group, Cambridge University Engineering Department

January 28, 2016

In this lecture...

Dialogue management architecture

Need for belief tracking

Generative approaches to belief tracking

Hidden Information State (HIS) dialogue model

Bayesian Update of Dialogue State (BUDS) model

Example dialogue

Hello, how may I help you?

I'm looking for a Thai restaurant. **inform(type=restaurant, food=Thai)**

What part of town do you have in mind?

Something in the centre. **inform(area=centre)**

Bangkok city is a nice place, it is in the centre of town and it serves Thai food.

What's the address? **request(address)**

Bangkok city is a nice place, their address is 24 Green street.

Thank you, bye. **bye()**

Spoken dialogue systems architecture

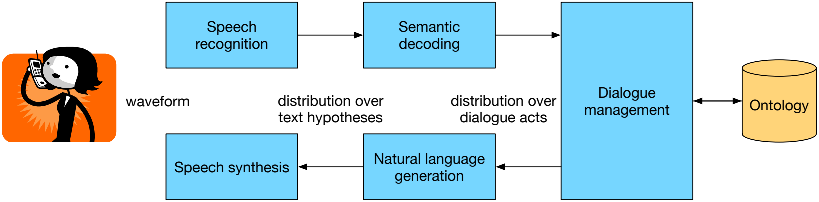

Dialogue management

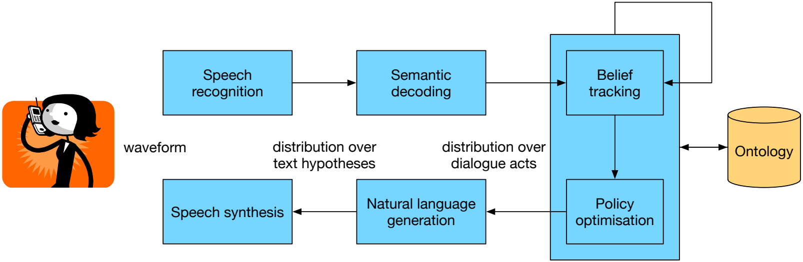

Example: 1-best input and no belief tracking

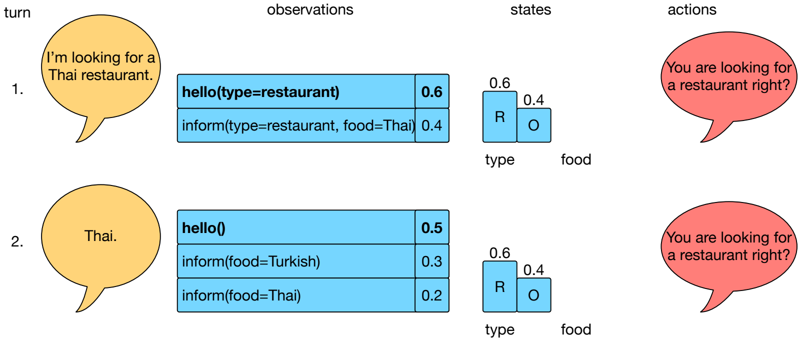

Example: N-best input and no belief tracking

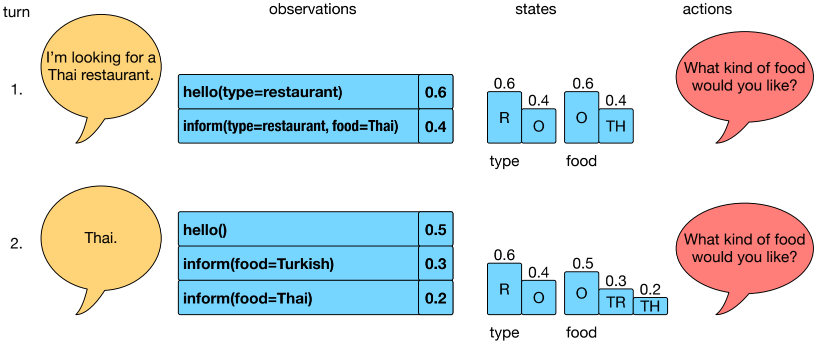

Example: N-best input with belief tracking

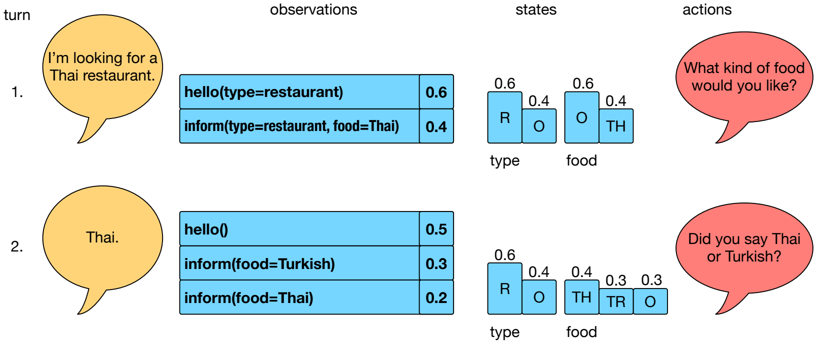

Elements of dialogue management

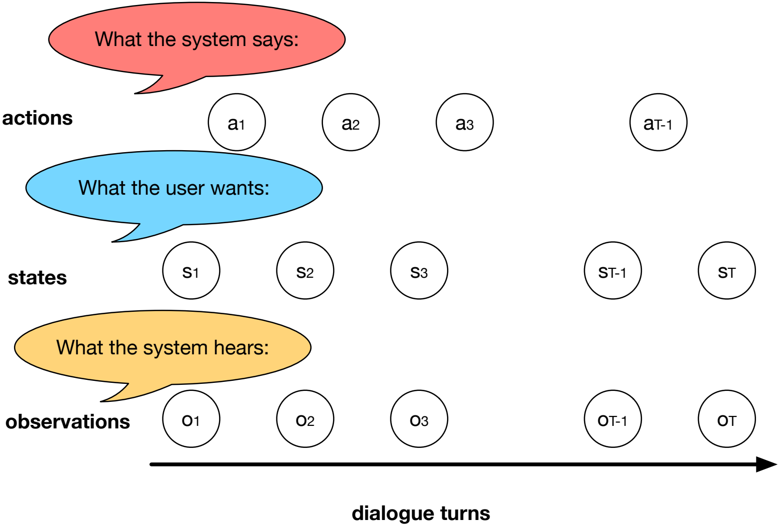

Challenges in dialogue modelling

- How to define the state space?
- How to tractably maintain the dialogue state?
- Which actions to take?

Solution: Define dialogue as a **control problem** where the behaviour can be automatically learned.

Dialogue management as Markov decision process

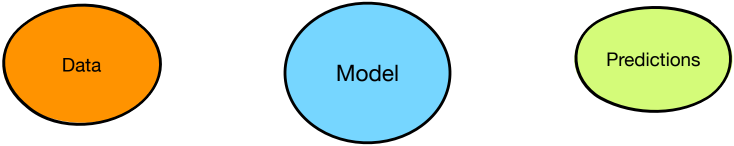

- Dialogue states
- Reward – a measure of dialogue quality

- Markov decision process

- Optimal system actions

Theory: Bayesian networks

- Bayesian network is a directed acyclic graph where nodes represent random variables and the arrows represent conditional independence assumption.
- Dynamic Bayesian network is a Bayesian network which repeats its structure at each point in time.

Theory: Markov decision process

s_t dialogue states

a_t system actions

r_t rewards

$p(s_{t+1}|s_t, a_t)$ transition probability

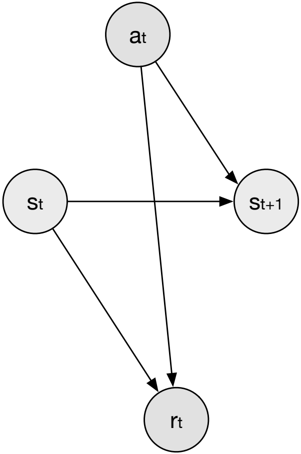

Dialogue as a Markov decision process?

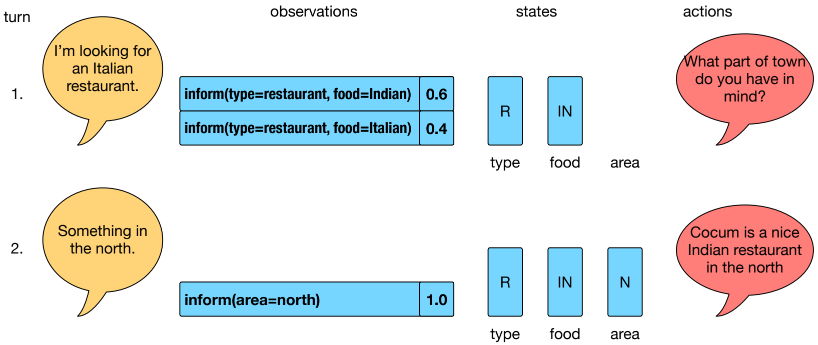

Dialogue management as partially observable Markov decision process

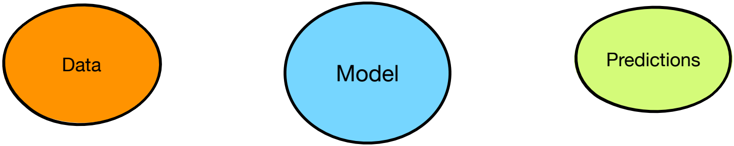

- Noisy observations of dialogue state
- Reward – a measure of dialogue quality

- Partially observable Markov decision process

- Distribution over possible dialogue states – **belief state**
- Optimal system actions

Generative vs discriminative models in belief tracking

Discriminative models: the state depends on the observation

$$b(s_t) = p(s_t|o_t)$$

Generative models: the state generates the observation

$$b(s_t) = \frac{p(s_t, o_t)}{\sum_{s_t} p(s_t, o_t)} \propto p(o_t|s_t)p(s_t)$$

Partially observable Markov decision process

- State generates a noisy observation $p(o_t|s_t)$ – the **observation probability**

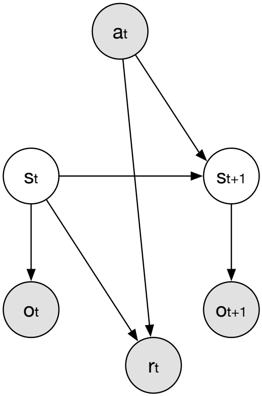

- State is unobservable and depends on the previous state and the action: $p(s_{t+1}|s_t, a_t)$ - the **transition probability**

Theory: Belief propagation

Probabilities conditional on the observations

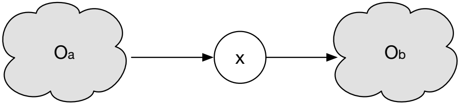

Interested in marginal probabilities $p(x|O)$, $O = O_a \cup O_b$

$$p(x|O_b, O_a) \propto p(x, O_b|O_a) = p(O_b|x, O_a)p(x|O_a) = p(O_b|x)p(x|O_a)$$

Theory: Belief propagation

Split O_b further into O_c and O_d

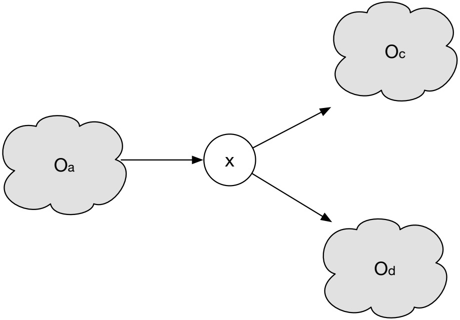

$$p(x|O_a, O_c, O_d) \propto p(O_c, O_d|x)p(x|O_a) = p(O_c|x)p(O_d|x)p(x|O_a)$$

Theory: Belief propagation

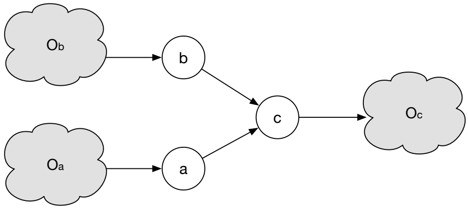

$$p(c|O_a, O_b) = \sum_{a,b} p(a|O_a)p(b|O_b)p(c|a,b)$$

$$p(O_c, O_b|a) \propto \sum_{b,c} p(O_c|c)p(b|O_b)p(c|a,b)$$

Theory: Belief propagation

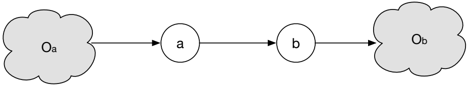

$$p(b|O_a) = \sum_a p(a|O_a)p(b|a)$$

$$p(O_b|a) = \sum_b p(O_b|b)p(b|a)$$

Belief state tracking

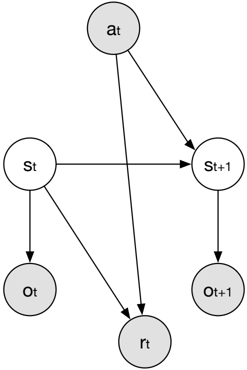

$$b(s_{t+1}) \propto p(o_{t+1}|s_{t+1}) \sum_{s_t} p(s_{t+1}|a_t, s_t) b(s_t)$$

Requires summation over all possible states at every dialogue turn – **intractable!**

Practical examples of POMDP systems

- POMDPs are normally intractable for everything but very simple cases
- However there are approximations which enable their use for real-world dialogue domains

  Hidden Information State (HIS) system [Young et al., 2010]
  Bayesian Update of Dialogue State (BUDS) system [Thomson and Young, 2010]

Requirements for belief tracking

Dialogue history: The system needs to keep track of what happened so far in the dialogue. This is normally done via the **Markov property**.

Task-orientated dialogue: The system needs to know what the user wants. This is modelled via the **user goal**.

Robustness to errors: The system needs to know what the user says. This is modelled via the **user act**.

Dialogue state factorisation

Decompose dialogue state into conditionally independent elements

user goal g_t

user action u_t

dialogue history d_t

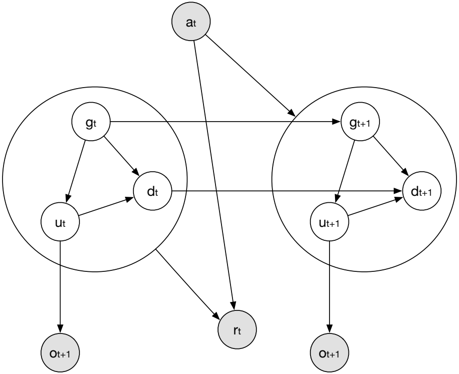

Belief update

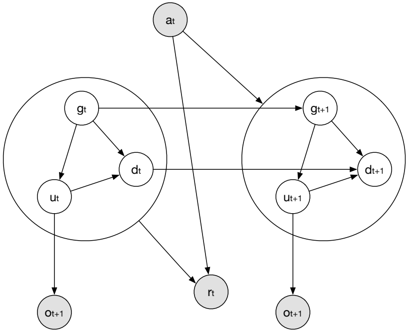

$$
\begin{aligned}
b(g_{t+1}, u_{t+1}, d_{t+1}) = \\
& p(o_{t+1}|u_{t+1}) \cdot \\
& p(u_{t+1}|g_{t+1}, a_t) \cdot \\
& \sum_{g_t} p(g_{t+1}|a_t, g_t) \cdot \\
& \sum_{d_t, u_t} p(d_{t+1}|d_t, g_{t+1}, u_{t+1}, a_t) \cdot \\
& b(g_t, u_t, d_t)
\end{aligned}
$$

- Requires summation over all possible goals – **intractable!**
- Requires summation over all possible histories and user actions – **intractable!**

Hidden Information State (HIS) dialogue state

Observation:
N-best list of user acts

User Goal:
Partitions of the goal space built according to ontology

Dialogue history:
Grounding states

Hypotheses:
Every combination of user act, partition and history

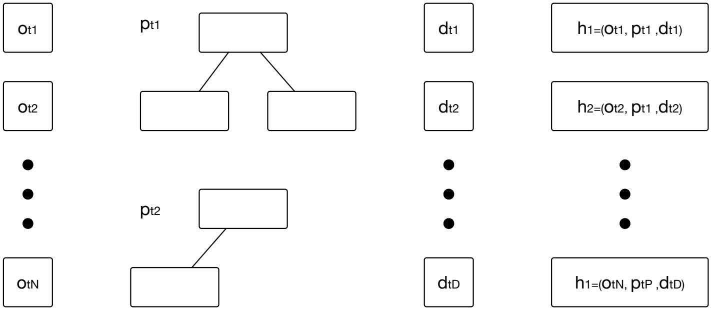

Belief state: Distribution over most likely hypotheses

HIS partitions

System: How may I help you?
request(task)

User: I'd like a restaurant in the centre.
inform(entity=venue,type=restaurant, area centre)

entity=venue

area=centre

type=restaurant

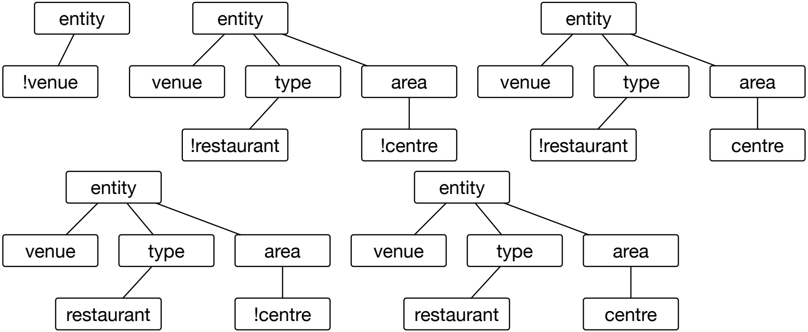

Pruning

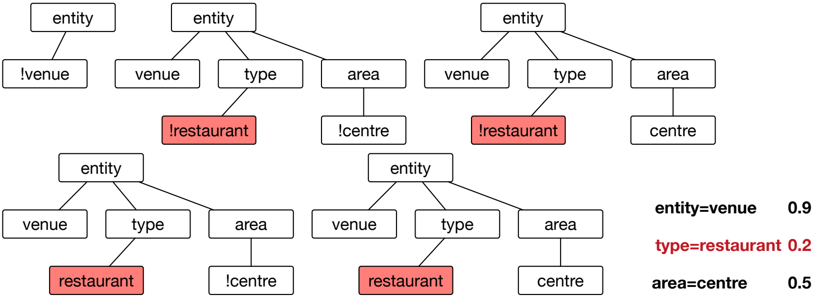
entity
!venue
entity
venue
type
area
!restaurant
!centre
entity
venue
type
area
!restaurant
centre
entity
venue
type
area
restaurant
!centre
entity
venue
type
area
restaurant
centre
entity=venue 0.9
type=restaurant 0.2
area=centre 0.5

Pruning

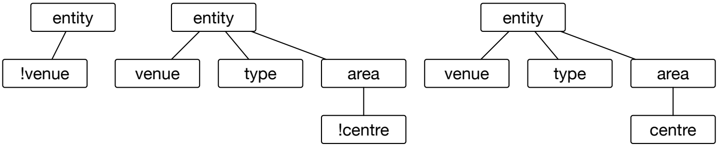

entity=venue **0.9**

area=centre **0.5**

Bayesian update of dialogue state model

- Further decomposes the dialogue state
- Produces tractable belief state update
- Transition and observation probability distributions can be parametrised and their shape learned

Bayesian network in the BUDS model

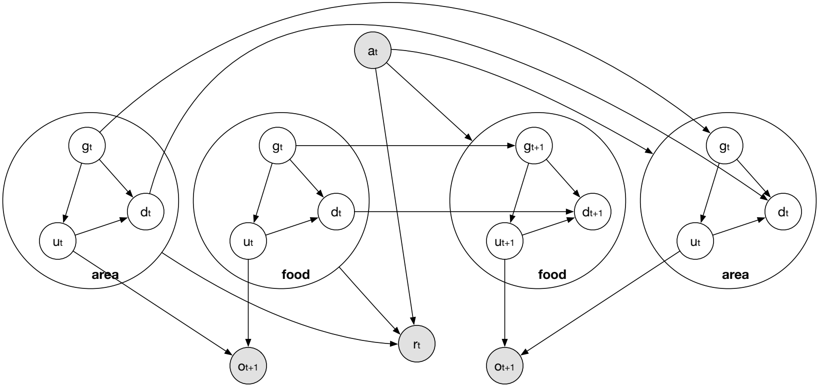

Belief tracking in the BUDS model

For each node x

- Start on one side and keep getting $p(x|O_a)$
- Then start on the other side and keep getting $p(O_b|x)$
- To get a marginal simply multiply these

Simple example

$p(o\|g)$	o :Thai	o :Turk.
g :Thai	0.8	0.2
g :Turk.	0.2	0.8

$p(g'\|g)$	g' :Thai	g' :Turk.
g :Thai	0.9	0.1
g :Turk.	0.1	0.9

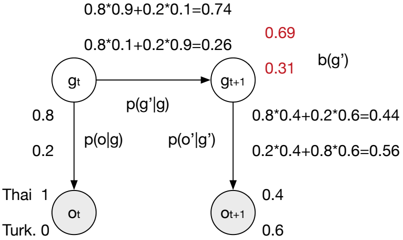

Learning of the shape of distributions

Expectation propagation

- Allows parameter tying
- Handles factorised hidden variables
- Handles large state spaces
- Does not require annotations but uses the output of the semantic decoder

Summary

- Properties of belief tracking for dialogue management include Markov assumption, being able to model the user goal and being robust to speech recognition errors
- Generative models for belief tracking are based on partially observable Markov decision processes
- Hidden Information State (HIS) model decomposes the dialogue state into the user goal, the user action and the dialogue history. Transitions are hand-crafted and the goals are grouped together to allow tractable belief tracking
- Bayesian Update of Dialogue State (BUDS) model further factorises the state which allows tractable belief tracking but also learning of the shapers of distributions via Expectation propagation

References

Thomson, B. and Young, S. (2010).
Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems.
Computer Speech and Language, 24(4):562–588.

Young, S., Gašić, M., Keizer, S., Mairesse, F., Schatzmann, J., Thomson, B., and Yu, K. (2010).
The Hidden Information State model: A practical framework for POMDP-based spoken dialogue management.
Computer Speech and Language, 24(2):150–174.