Adaptive Multimodal Dialogue Management based on the Information State Update Approach

Kallirroi Georgila and Oliver Lemon

University of Edinburgh

TALK project

Overall Coordination:

Scientific Coordination:

TALK general aim

- The project will generalise the Information State Update approach to dialogue management, as developed in the TRINDI (Larsson and Traum, 2000) and SIRIDUS (Lewin et al., 2000) projects, in order to develop adaptive multimodal dialogue systems.

TALK research themes

- Unifying multimodality and multilinguality
- Automatic generation and reconfiguration of multimodal interfaces
- Multimodal presentation in the Information State Update approach
- Learning and adaptivity
  - Reinforcement Learning for dialogue management
  - Complex dialogue states from the Information State Update approach

Information State Update approach

- The Information State Update (ISU) approach allows a declarative representation of dialogue modelling.
- "The term Information State of a dialogue represents the information necessary to distinguish it from other dialogues, representing the cumulative additions from previous actions in the dialogue, and motivating future action" (Larsson and Traum, 2000).

Example information state

```
lastspeaker: user
turn: system
output:
  < hello, welcome to the edinburgh informatics
    automatic information system. how may i help you? >
input: < i would like information about restaurants >
lastmoves:
  < [i would like information about restaurants],u,
     [([greet],s),([ask_how_to_help],s)] >
filledslotsvalues:
  < [([ask_how_to_help],s)],[[restaurants]] >
oplansteps:
  ( [ask_user_restaurant_type] , [release_turn] )
nextmoves: < [ask_user_restaurant_type],s >
int: < [release_turn] >
. . .
```

Dialogue strategy

A dialogue strategy would be for example for the system to decide on:

- the type of confirmation
  - explicit (“Are you leaving from Edinburgh?”)
  - implicit (“Leaving from Edinburgh, where would you like to fly?”)
  - none
- the modality it would use to present the requested information
  - speech
  - text
  - icons

Reinforcement Learning (RL)

- Dialogue is modelled as a Markov Decision Process (MDP) (Levin and Pieraccini, 1997)
- Choose the action *a* which maximizes the expected reward *Q(s,a)* given the state *s*

$$Q(s,a) = R(s,a) + \sum_{s'} P(s' \mid s,a) \max_{a'}(Q(s',a'))$$

- Estimate *P(s′ |s,a)* from users' behavior
- Estimate *Q(s,a)* iteratively from sample dialogues

Information State Update approach with policy learning

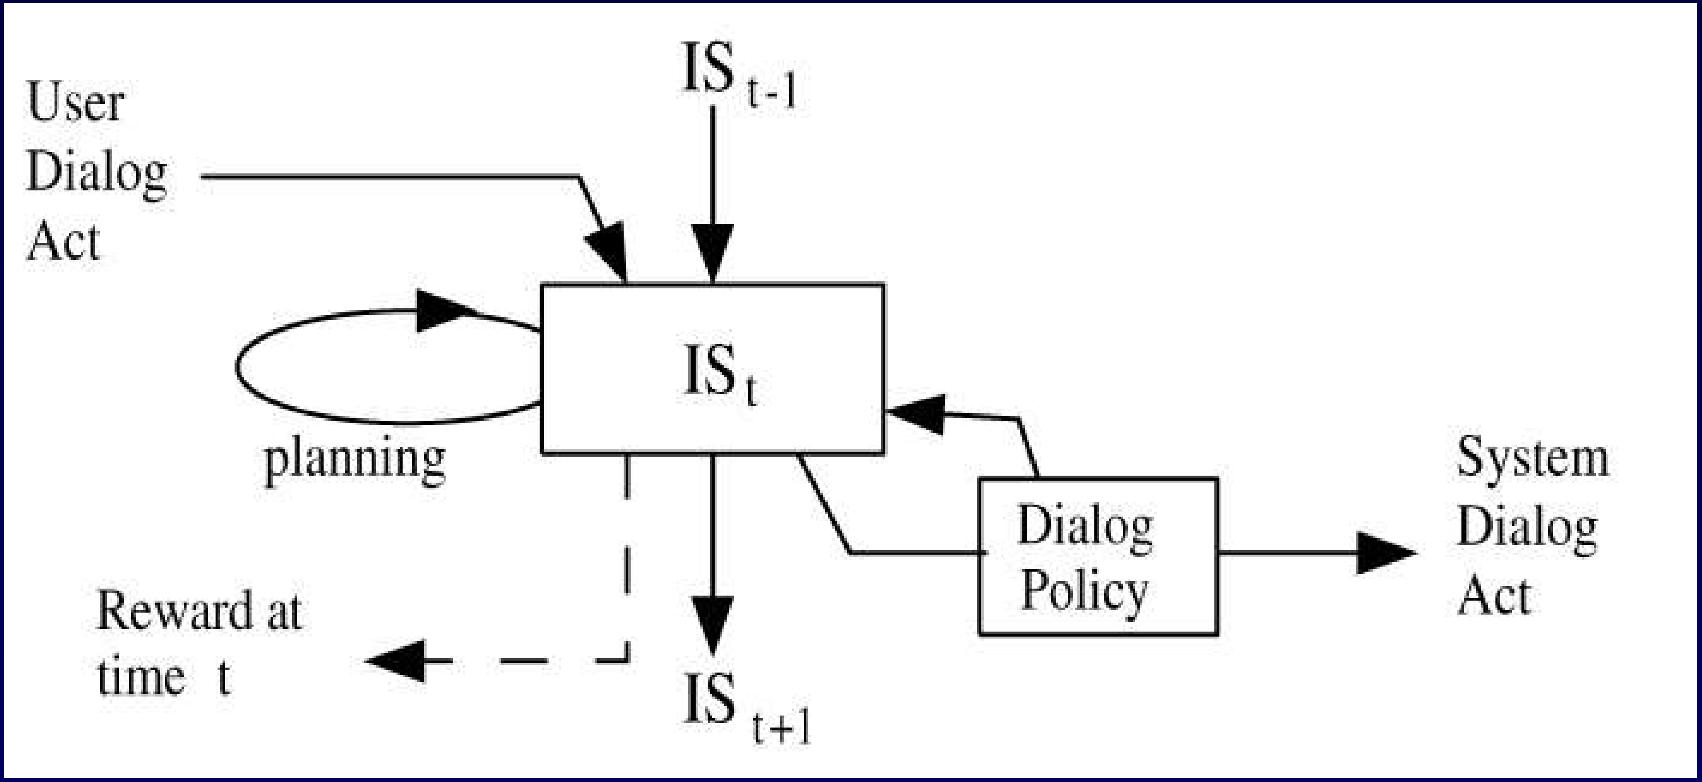

Possible sources of data for learning

- Real human-machine interactions (through an ASR system)
- Large amounts of corpus data
- Simulated human-machine interactions (virtual user) (Scheffler and Young, 2000-2002)

TALK baseline system

- DIPPER (Bos et al., 2003) for dialogue management
- ATK (Young, 2004) for speech recognition
- Festival (Taylor et al., 1998) for speech synthesis
- O-Plan (Currie and Tate, 1991) for dialogue planning and content planning and structuring

Example information state definition

```
infostate(record([is:record([
            lastspeaker:atomic,
            turn:atomic,
            input:stack(atomic),
            lastinput:stack(atomic),
            output:stack(atomic),
            nextmoves:stack(Acts),
            lastmoves:stack(Acts),
            filledslotsvalues:stack(atomic),
            filledslots:stack(atomic),
            int:stack(Acts)])])) :-
            Acts = record([pred:atomic,
                           dp:atomic,
                     prop:record([pred:atomic,
                     args:stack(atomic)])]).
```

Example DIPPER update rule

```
urule(generation,
    [;;; CONDITIONS:
     top(is^int)=[release_turn],
     is^lastspeaker=user,
     prolog(checkfilledslots(top(is^nextmoves),
                  is^filledslots,Z)),
     Z=0,
    ],
    [;;; EFFECTS:
     prolog(reverse_and_utter(is^nextmoves,
            X,Y)),
     push(is^lastmoves,X),
     clear(is^output),
     push(is^output,Y),
     solve2(callfestival(Y,_X)),
     assign(is^lastspeaker,system),
     assign(is^turn,user)
    ] ).
```

The Graphical User Interface of DIPPER

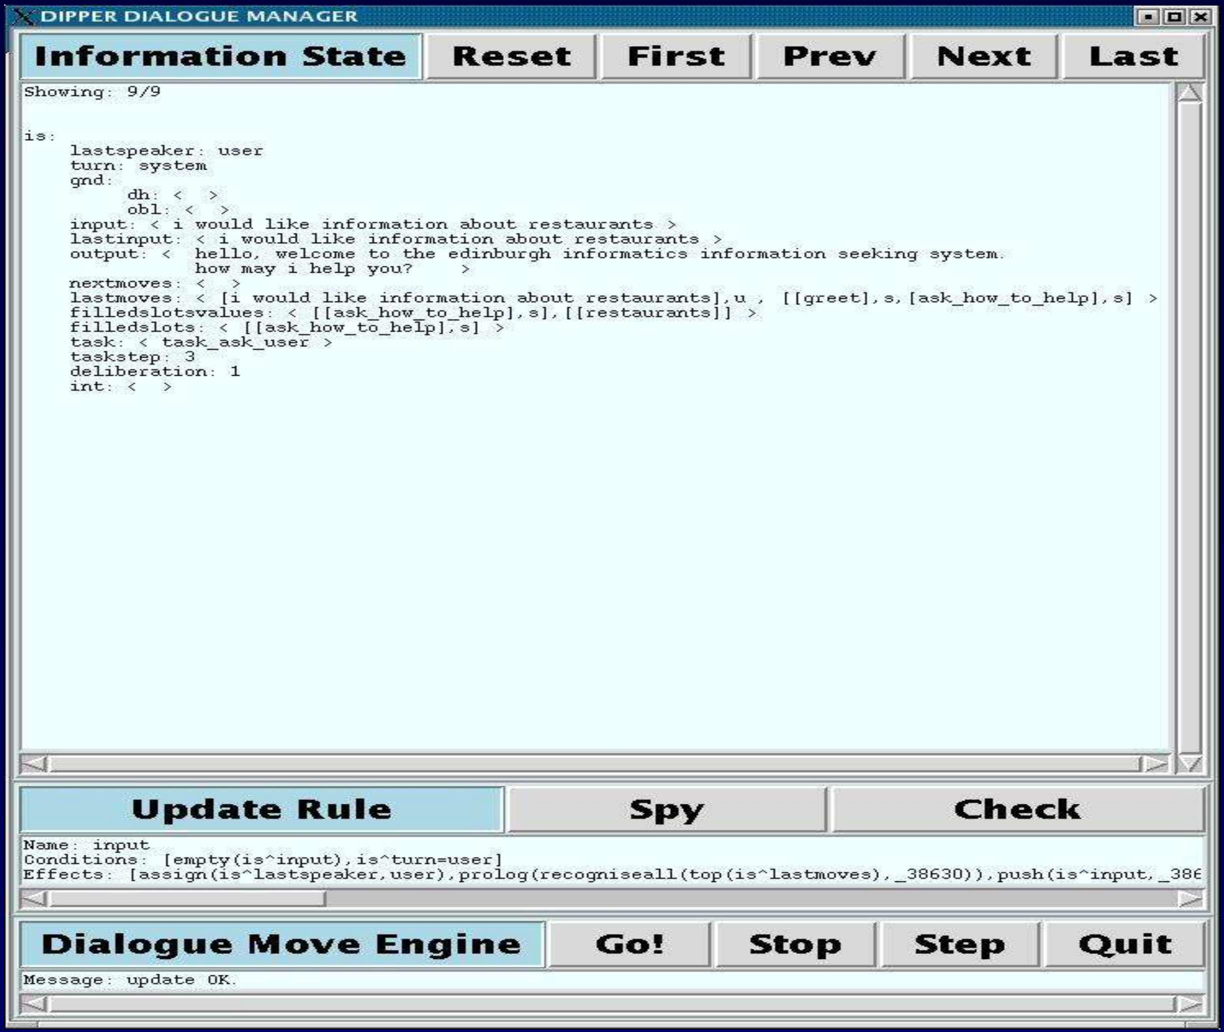

Communicator 2000 corpus

- Flight information, car rental, hotel booking
- 662 human-machine dialogues
- 9 different travel planning systems
- 60-79 dialogues per system
- Transcription of user input
- Only system utterances are tagged

(Walker et al., 2001)

Example Communicator data

SYS: Welcome.

SYS: You are logged in as a guest user of Ay T and T Communicator. You may say repeat, help me out, start over, or, that’s wrong, you can also correct and interrupt the system at any time.

SYS: What airport woodja like to fly out of?

USER: ASR: <CITY>HONOLULU HAWAII</CITY>
TRANS: <CITY>HONOLULU HAWAII</CITY>

SYS: Leaving from <CITY>Honolulu</CITY>,

SYS: And, what city are you flying to?

USER: ASR: <CITY>DALLAS TEXAS</CITY>
TRANS: <CITY>DALLAS TEXAS</CITY>

SYS: Flying from <CITY>Honolulu</CITY> to <CITY>Dallas Fort Worth</CITY>,

SYS: What date would you like to fly?

USER: ASR: <DATE_TIME>WEDNESDAY NOVEMBER ELEVENTH</DATE_TIME>
TRANS: <DATE_TIME>WEDNESDAY NOVEMBER ONE</DATE_TIME>

. . .

Initial data collection
Cambridge SACTI-1 corpus

- SACTI stands for Simulated ASR-Channel: Tourist Information (Stuttle et al., 2004, Williams and Young, 2004).
- Tourist information, with route descriptions
- Human-human data
- On-line transcription of user input
- Speech recognition error simulation
- In a new data collection (not part of SACTI-1 corpus) highlighting and clicking on maps is also included

Example SACTI-1 data

hello how can i help

AH I'M LOOKING FOR A GOOD RESTAURANT IN THE TOWN

right there's a number of restaurants in town %um what sort of food are you looking to -- to eat

I'M LOOKING FOR A RESTAURANT NEAR THE CINEMA

okay there's a restaurant very near the cinema it's a -- a relaxed chinese restaurant called noble nest

NOBLE NEST AND AH WHERE IS IT EXACTLY

it's on the corner of north road and fountain road

AND AH WHAT IS THE PRICE OF FOOD THERE

%er the food there is %er fourteen pounds per person

AH IT'S A CHINESE RESTAURANT RIGHT

that's right yes

AND HOW TO REACH NOBLE NEST FROM HOTEL ROYAL

right okay from the hotel royal it would probably be best to catch the bus outside the hotel royal %um which will take you -- probably catch the bus to art square and then walk %um from art square that being the closest bus stop

. . .

Initial objectives regarding RL

- Which aspects of dialogue management are amenable to learning and what reward functions are needed for these aspects?
- What representation of the dialogue state best serves this learning?
- What Reinforcement Learning methods are tractable with large scale dialogue systems?

Previous applications of RL to dialogue management

- (Levin, Pieraccini and Eckert, 2000), (Singh, Litman, Kearns and Walker, 2002)
- Choose between a small number of actions
  - Initiative: system / user / mixed
  - Confirmation: explicit / none
- Have a small number of possible states
- Use RL methods which would not scale up to large action sets and large state spaces

Challenges regarding RL

- Tractable Reinforcement Learning with complex actions and large numbers of state features
- Learning generic strategies which can be applied to many domains
- Discovering useful features of the dialogue history
- Including partially observable features of the state (using POMDP models)

Summary

- Adaptation and learning of multimodal dialogue strategies is an important theme in the TALK project
- TALK uses the Information State Update approach, with large state representations
- Reinforcement Learning will be used to learn dialogue management
- Challenges in tractable learning with these large complex representations

For on-going project info visit the TALK website

http://www.talk-project.org