Slides for a tutorial presented at the
Natural Language Generation, Summarisation, and Dialogue Systems Summer School
July 20-24, 2015, Aberdeen.

The Open University

Introduction to Dialogue Systems

Paul Piwek
Centre for Research in Computing

Purpose

- Have an understanding of the key ideas behind most dialogue systems.
- Be familiar with some current research challenges.

Purpose

Variety

- Chatbots (Eliza, ...)
- Travel information systems (Philips Timetable system, ...)
- Intelligent assistants (Siri, Google Now, Cortana, IKEA's Anna, ...)
- Collaborative planners (John and Mary, TRAINS, ...)
- ...

Outline

- Dialogue games
- Systems
  - Reactive
  - Agenda-driven
  - *Evaluation*
  - *Machine learning*
- Non-cooperation
- Incrementality

Dialogue Games

Wittgenstein (1958: 77)

Imagine this language:–

1). Its function is the communication between a builder A and his man B. B has to reach A building stones. There are cubes, bricks, slabs, beams and columns. The language consists of the words “cube”, “brick”, “slab”, “column”. A calls out one of these words upon which B brings a stone of a certain shape. Let us imagine a society in which this is the only system of language. The child learns this language from the grown-ups by being trained to its use.

Dialogue Games

Bunt and Van Katwijk (1979: 266-268)

“**Dialogue acts as elements of a language game**”

“What does it mean to view something as a game? A game is an activity in which the participants take turns in performing certain actions, chosen from the set of ‘legitimate moves’, in order to arrive at a preferred situation (‘favourable position’). Comparing this characterisation of a game with the characterisation of informative dialogues […] we can indeed view [dialogue] as a game, sequences of dialogue acts corresponding to moves, and the position that the players want to reach being a desired state of knowledge (…) think of a ‘position’ as an independent concept, as a ‘configuration of pieces’, as is for instance common in chess.”

Dialogue Games

Definition Beun (2001), Piwek (1998), Larsson & Traum (2003), Bos et al. (2003), Ginzburg (2012), ...

A dialogue game consists of two principal components:

*A **dialogue store***, for keeping track of the current position.

Dialogue rules which specify, for any given point in a dialogue, which *dialogue acts* are permitted at that point in the dialogue and how the store changes as a result of those actions. They are divided into two types of rules:

a) ***update rules***, which specify how the dialogue store evolves in the course of a dialogue.

b) ***generation rules***, which specify which dialogue acts are legitimate given a specific position (as recorded in the dialogue store).

Dialogue Games

Definition

Additionally, each dialogue participant needs a ***dialogue strategy***.

Given a set of available legitimate dialogue acts for a position, the strategy picks the act which will actually be played

Dialogue Games

A’s turn

Dialogue Store 1

Dialogue Games

A's turn

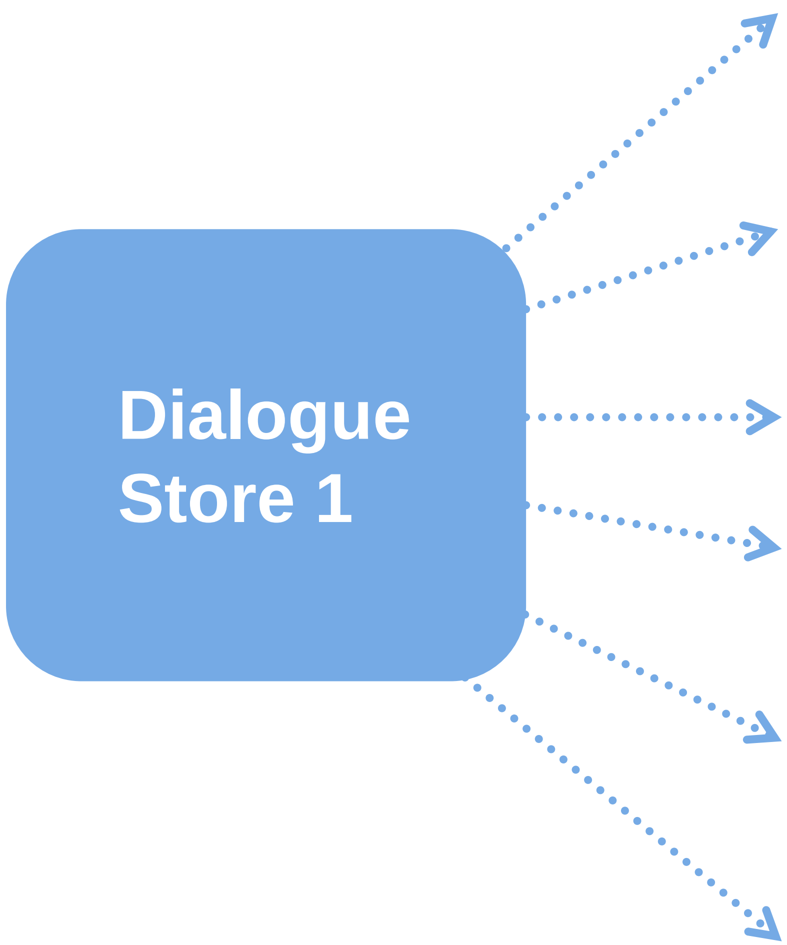

legitimate moves according to the generation rules

Dialogue Games

A’s turn

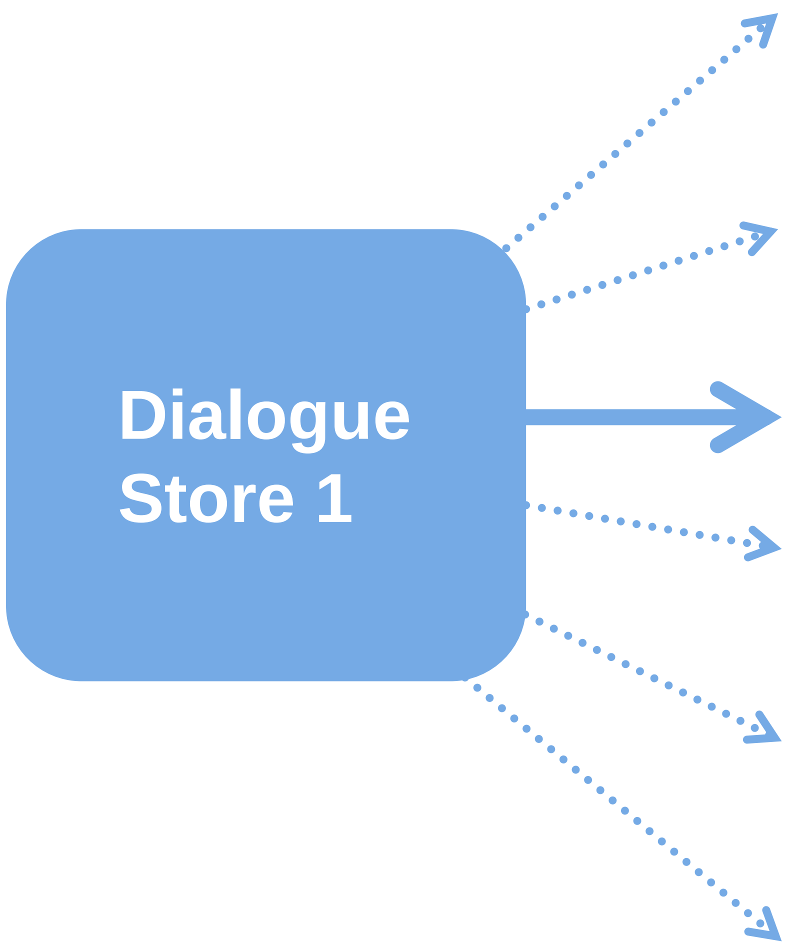

strategy selects
one specific move

Dialogue Games

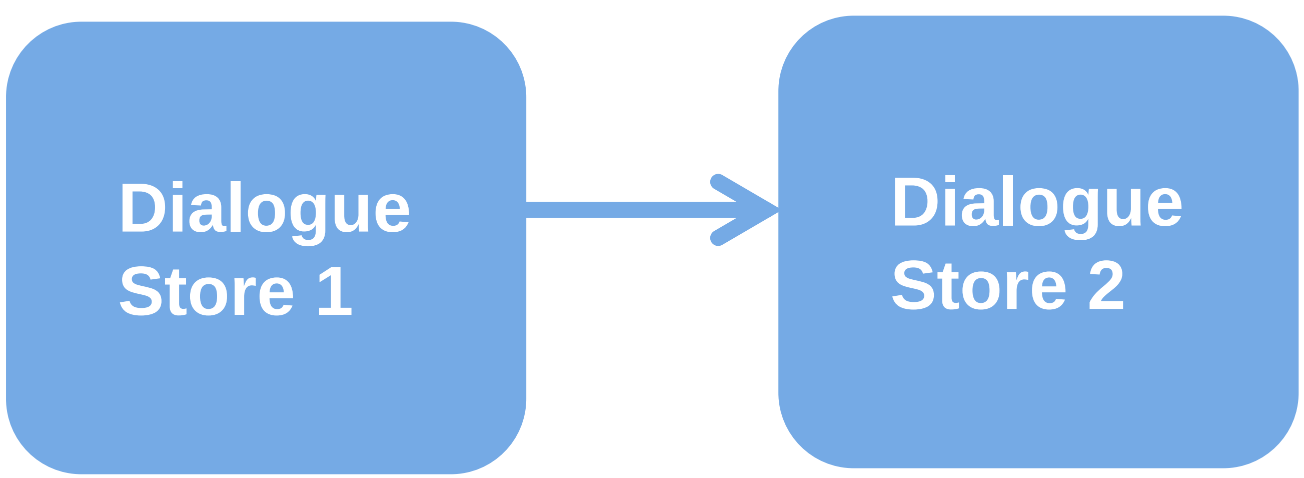

strategy selects
one specific move

Dialogue Games

B's turn

update rules

Dialogue Games

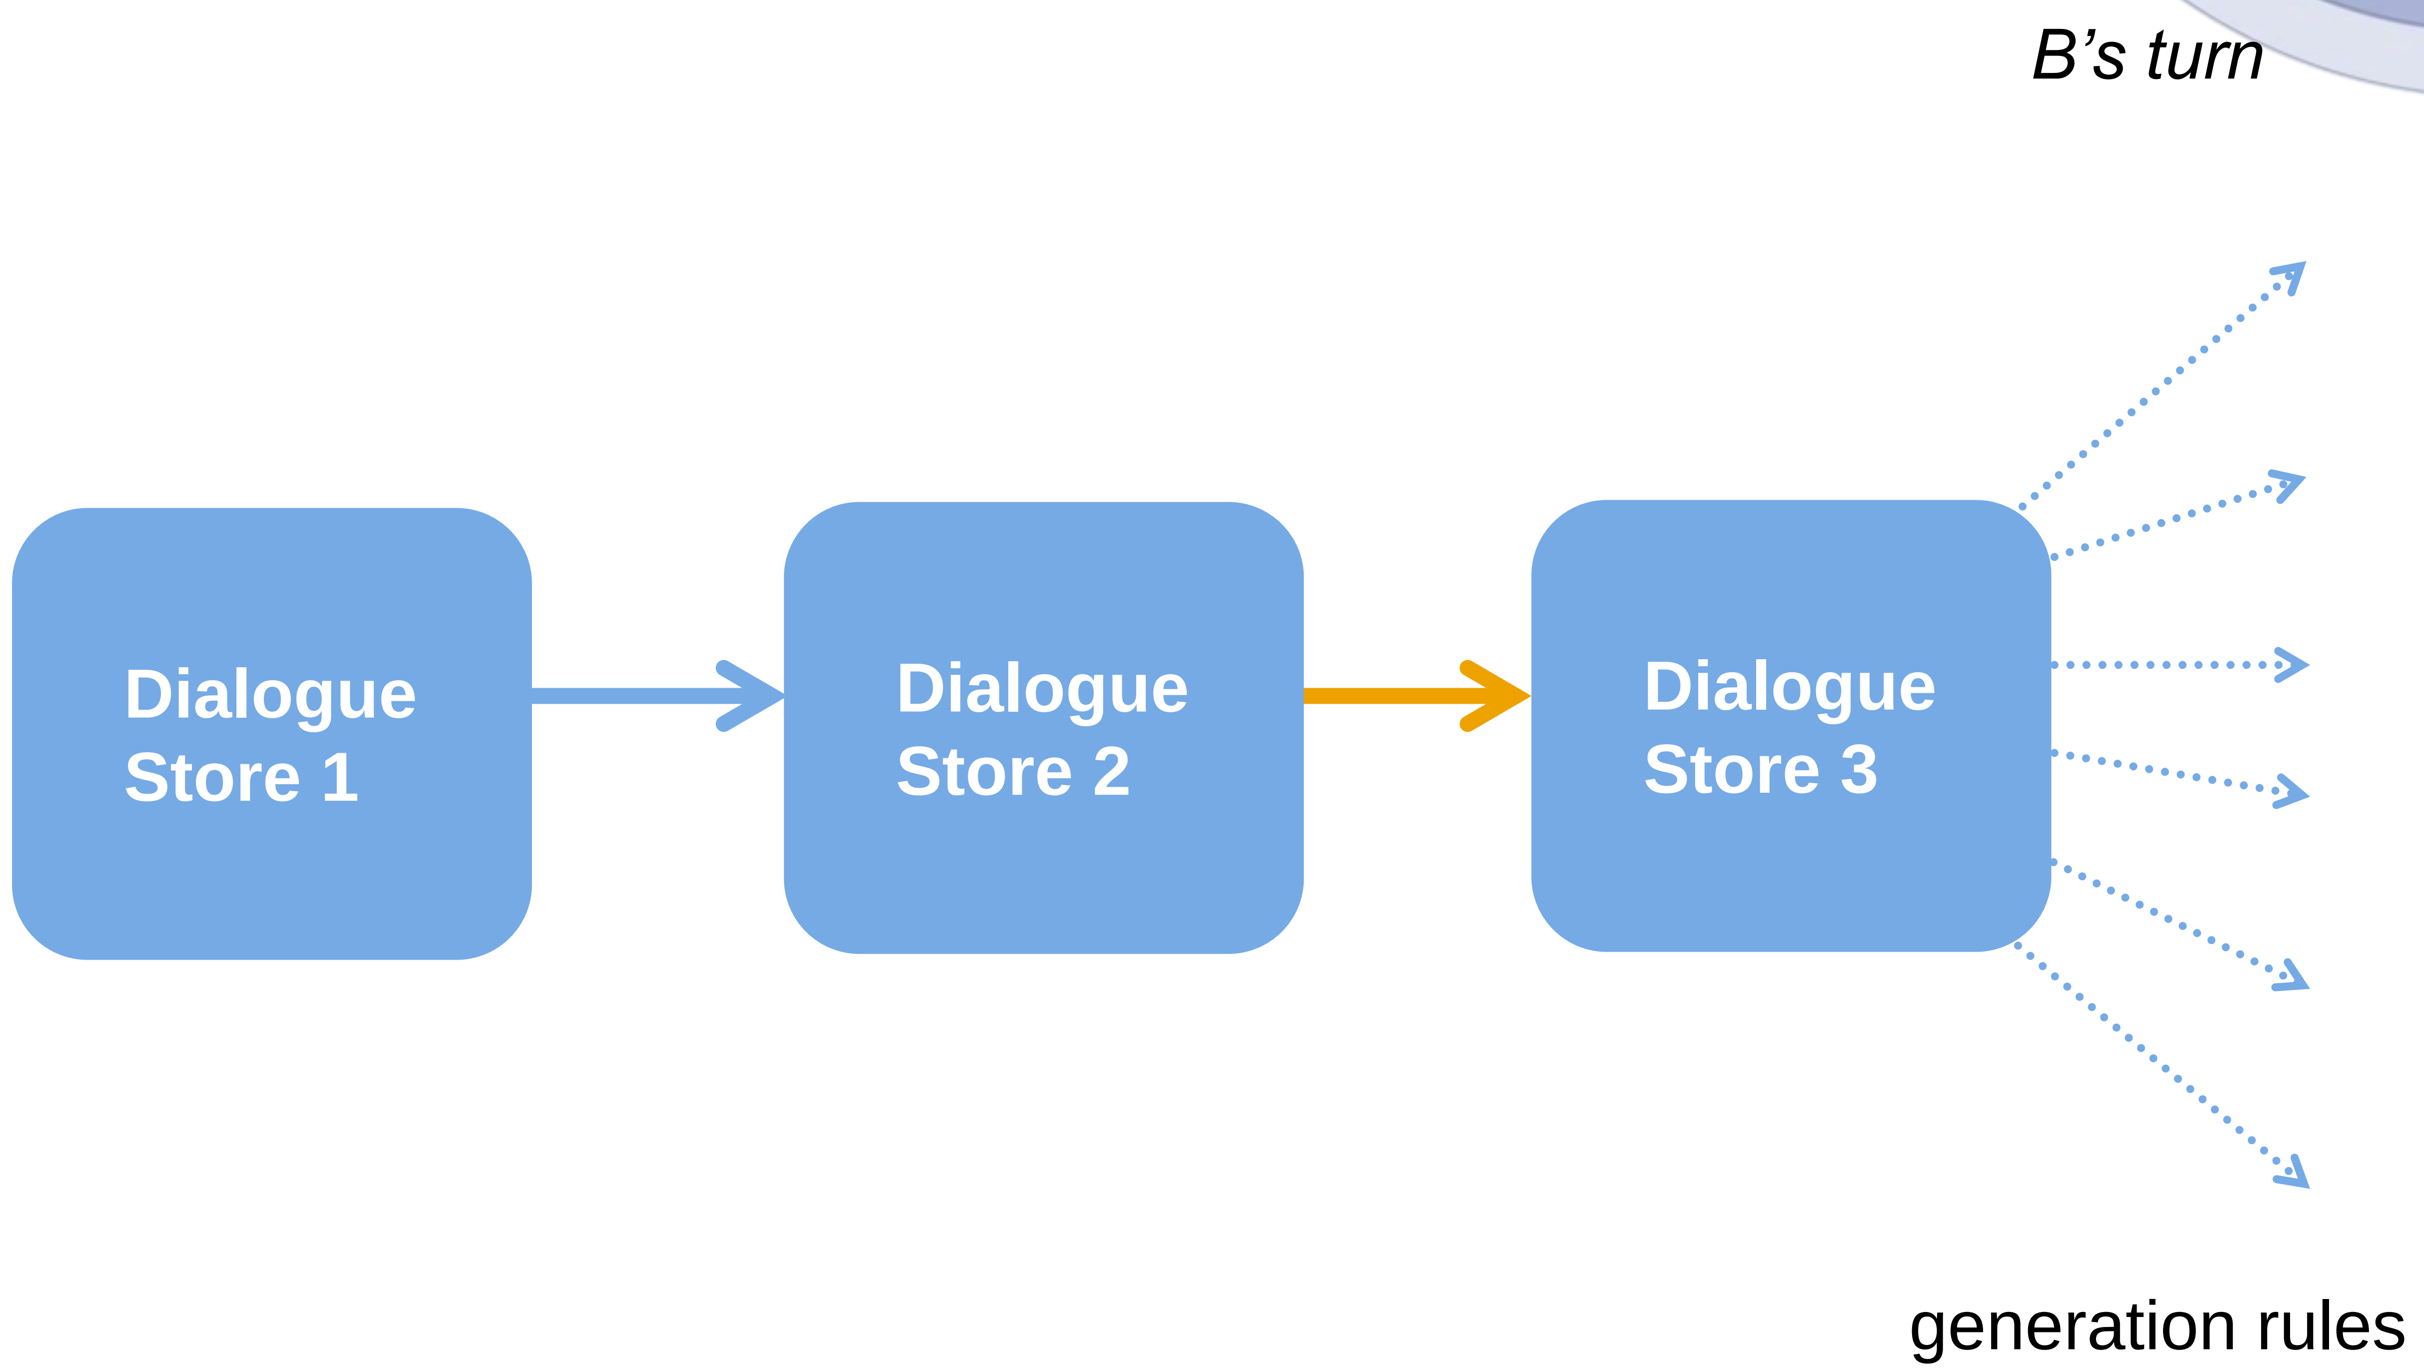

Dialogue Games

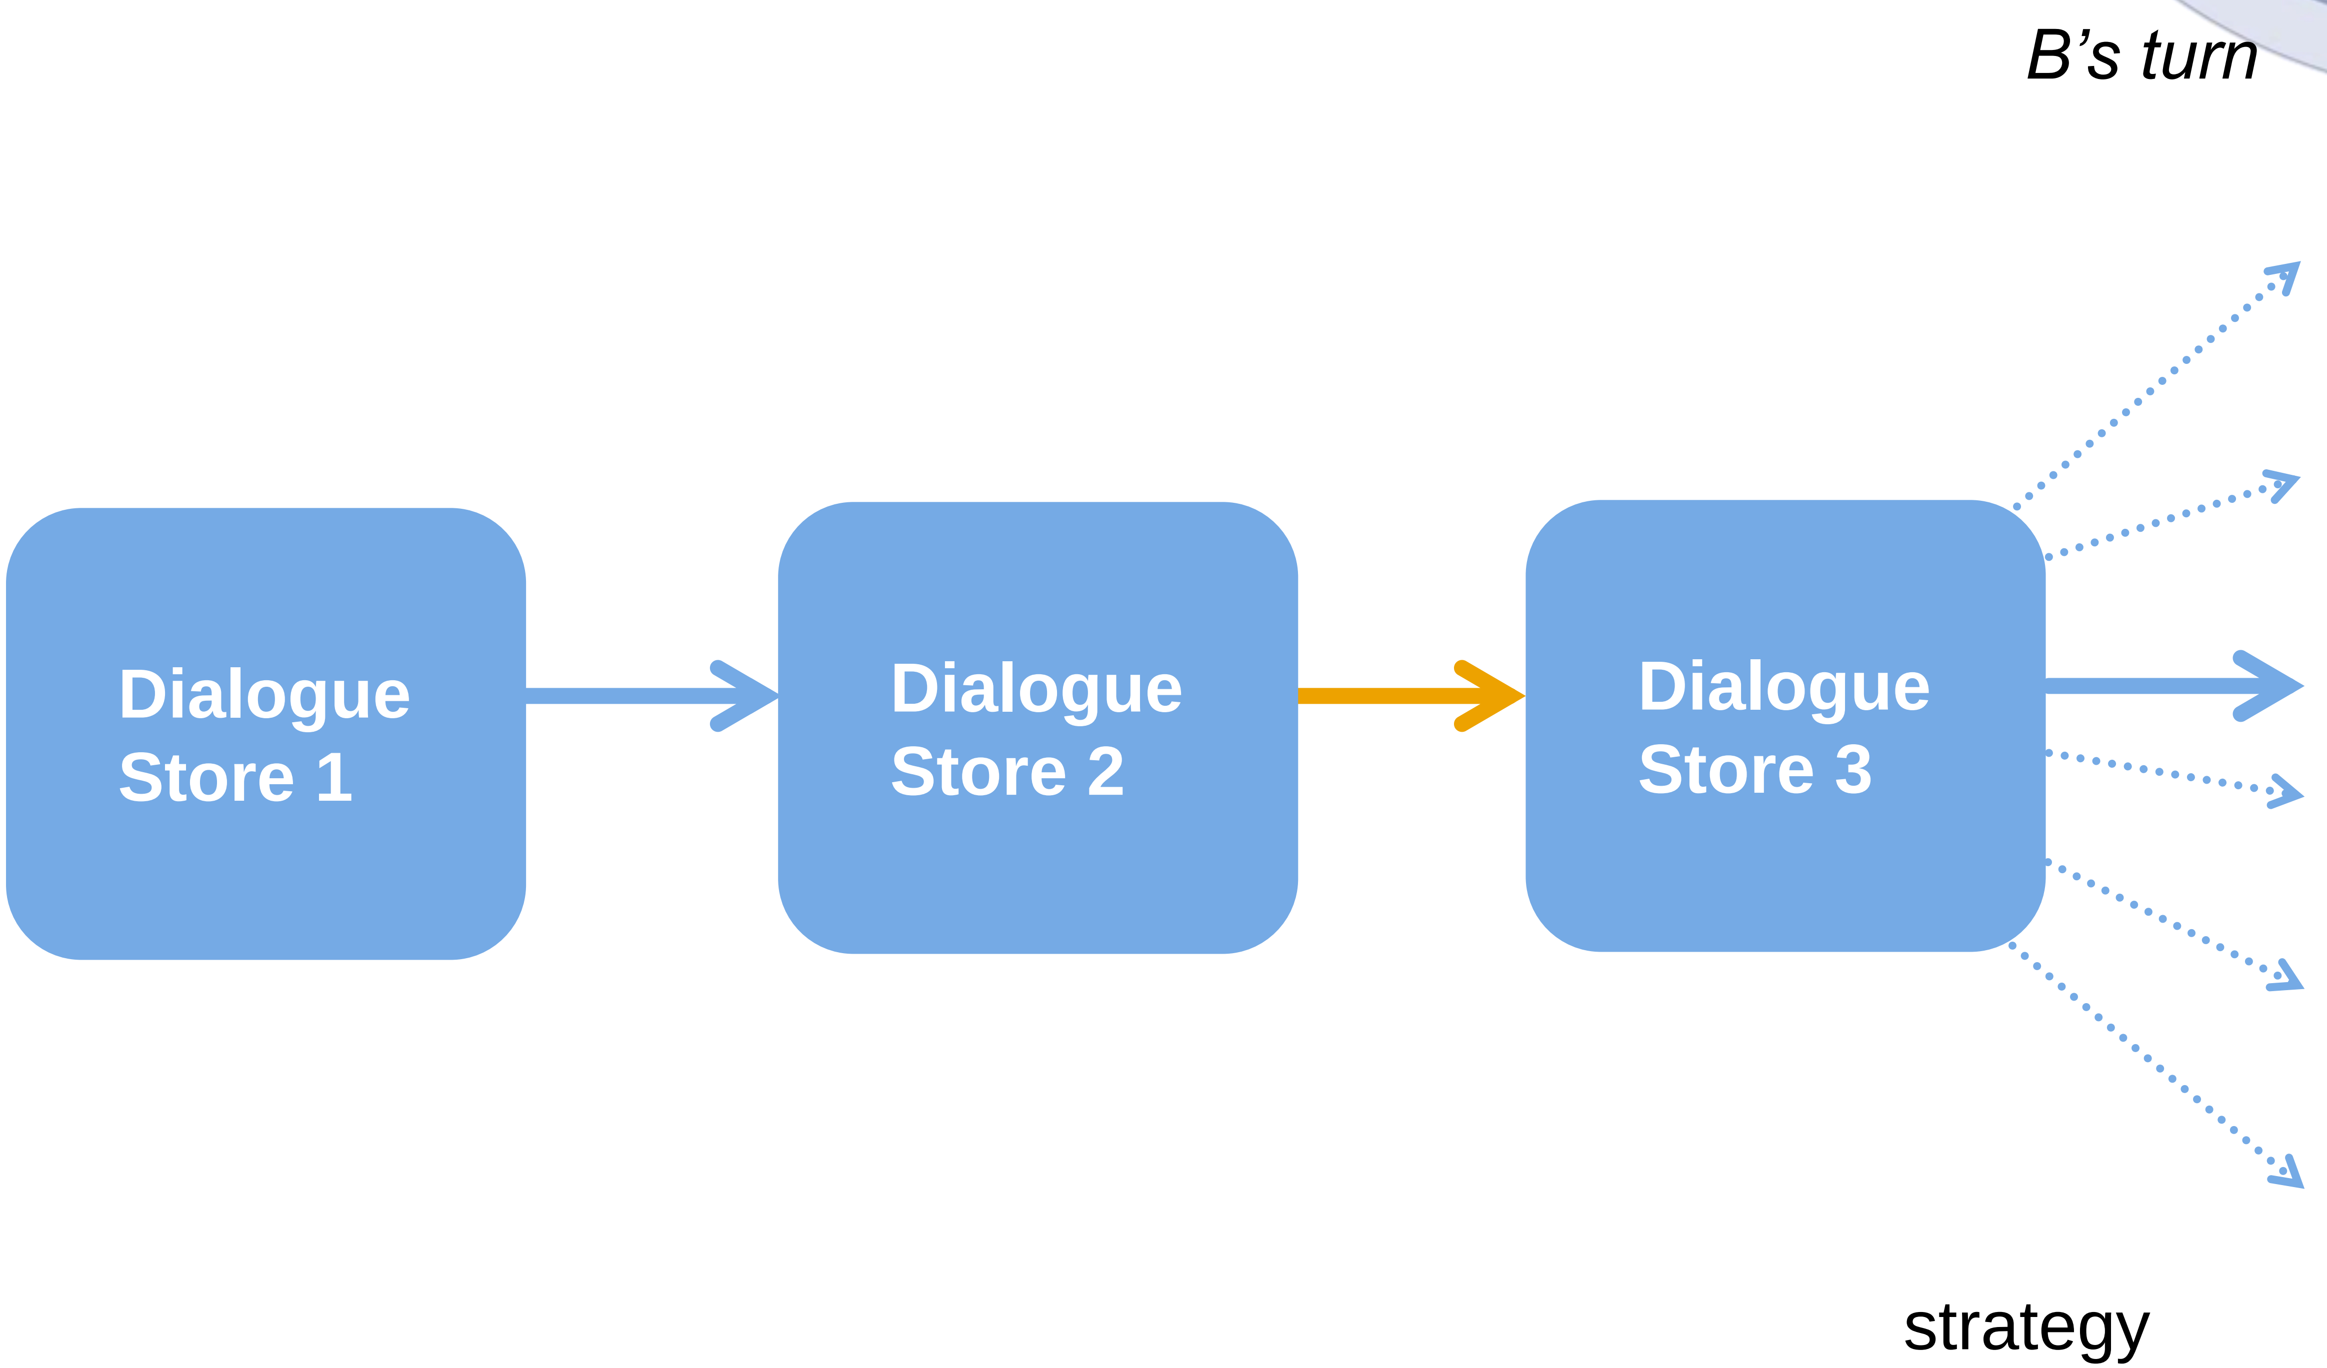

Reactive dialogue systems

ELIZA (Weizenbaum, 1966)

1	Human	Men are all alike.
2	ELIZA	in what way
3	Human	They are always bugging us about something or other.
4	ELIZA	can you think of a specific example
5	Human	Well, my boyfriend made me come here.
6	ELIZA	your boyfriend made you come here
7	Human	He says I'm depressed much of the time.
8	ELIZA	I am sorry to hear you are depressed
9	Human	It's true. I am unhappy.
10	ELIZA	do you think coming here will help you not be unhappy
11	Human	I need some help, that much seems certain.
12	Eliza	what would it mean to you if you got some help
13	Human	(…)

Reactive dialogue systems

ELIZA (Weizenbaum, 1966: 42-43)

The whole issue of the credibility (to humans) of machine output demands investigation. Important decisions increasingly tend to be made in response to computer output. (…) ELIZA shows, if nothing else, how easy it is to create and maintain the illusion of understanding, hence perhaps of judgement deserving credibility. A certain danger lurks there.’ (Weizenbaum, 1966: 42-43).

Reactive dialogue systems

ELIZA's game

- Store
- Rules for updates and generation

Reactive dialogue systems

ELIZA's game

Store - current input

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

He says I'm depressed much of the time.

Store - current input

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

He says I'm depressed much of the time.

Store - current input

Update rule: *Replace* "I'm" *with* "you are".

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

He says you are depressed much of the time.

Store - current input

Update rule: *Replace* "I'm" *with* "you are".

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

He says you are depressed much of the time.

Store - current input

decomposition template

Update rule: *Replace* "I'm" *with* "you are".

Generation rule: *Map* you are sad / unhappy / depressed / sick ... *to* I'm sorry to hear you are sad / unhappy / depressed / sick

reassembly pattern

Reactive dialogue systems

ELIZA's game

Human: He says I'm depressed much of the time.

I'm sorry to hear you are depressed

Store - current input

decomposition template

Update rule: *Replace* "I'm" *with* "you are".

Generation rule: *Map* you are sad / unhappy / depressed / sick ... *to* I'm sorry to hear you are sad / unhappy / depressed / sick

reassembly pattern

Reactive dialogue systems

ELIZA's game

But how can we have a set of generation rules that covers every possible input?

Generation rule: *Map to*

I see / that's interesting

reassembly pattern

Reactive dialogue systems

ELIZA's game

But how can we have a set of generation rules that covers every possible input?

input

memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why *Item or*
earlier you said *Item*

Reactive dialogue systems

ELIZA's game

Human: Well, my boyfriend made me come here.

input

memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce* let's discuss further why Item *or* earlier you said Item

Reactive dialogue systems

ELIZA's game

Human: Well, my boyfriend made me come here.

Well, my boyfriend made me come here.	input
	memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why Item *or*
earlier you said Item

Reactive dialogue systems

ELIZA's game

Human: Well, my boyfriend made me come here.

Your boyfriend made you come here.	input
	memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why Item *or*
earlier you said Item

Reactive dialogue systems

ELIZA's game

Human: Well, my boyfriend made me come here.

Your boyfriend made you come here.	input
Your boyfriend made you come here.	memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why Item *or*
earlier you said Item

Reactive dialogue systems

ELIZA's game

ELIZA: Your boyfriend made you come here.

input

Your boyfriend made you come here.

memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce* let's discuss further why Item *or* earlier you said Item

Reactive dialogue systems

ELIZA's game

ELIZA: What else comes to mind when you think of your father?
Human: Bullies.

Bullies — input

Your boyfriend made you come here. — memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why Item *or*
earlier you said Item

Reactive dialogue systems

ELIZA's game

ELIZA: What else comes to mind when you think of your father?
Human: Bullies.
ELIZA: Earlier you said your boyfriend made you come here

input

memory

Update rule: *Put content* "Your ..." *of input in memory*.

Generation rule: *If none of the generation rules with a non-empty decomposition template applies, retrieve Item from memory and produce*
let's discuss further why Item *or*
earlier you said Item

Reactive dialogue systems

Beyond ELIZA - From direct matches to similarity

- Leuski & Traum (2008): relax way in which the input needs to match with the decomposition pattern. From matching to similarity.
- Think of a QA system as having a large list of QA pairs. The task is to find a pair given the user's input. This is essentially an information retrieval task.

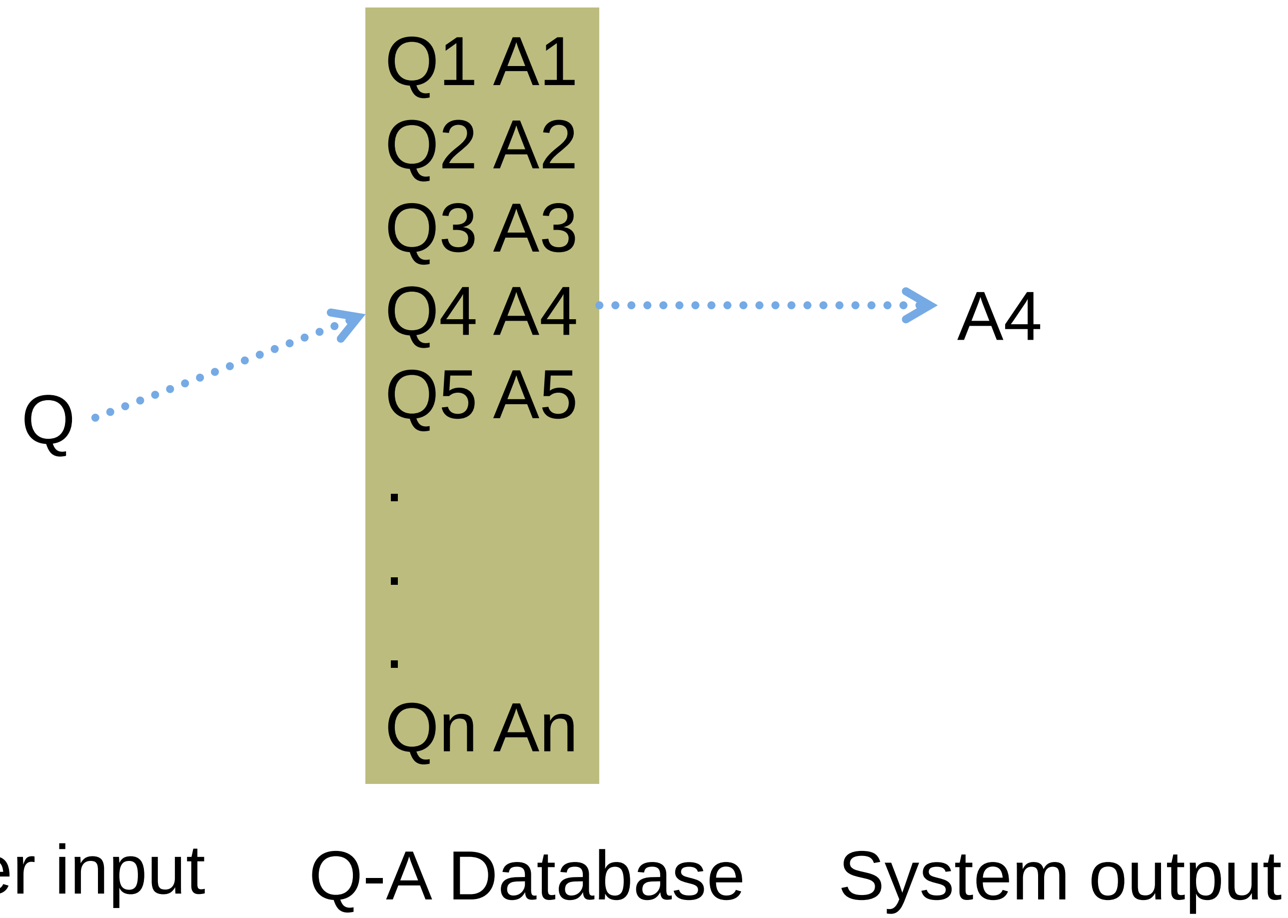

Reactive dialogue systems

Beyond ELIZA - Harvesting generation rules from text

CODA (Piwek & Stoyanchev, 2010; Kuyten et al., 2012)

1. Create a parallel monologue – dialogue corpus
2. Annotate the corpus with discourse relations and dialogue acts

Reactive dialogue systems

Beyond ELIZA - Harvesting generation rules from text

Monologue	**Dialogue** (from Twain 1919: 14 and 1)	
Text	*Speaker*	*Text*
One cannot doubt that he felt well.	OM	He felt well?
	YM	One cannot doubt it.
The metals are not suddenly deposited in the ores. It is the patient work of countless ages.	OM	Are the metals suddenly deposited in the ores?
	YM	No --
	YM	it is the patient work of countless ages.

Reactive dialogue systems

Beyond ELIZA - Harvesting generation rules from text

Monologue		**Dialogue** (from Twain 1919: 14 and 1)		
Text	*Rhetorical relation*	*Speaker*	*Text*	*Dialogue act*
One cannot doubt that he felt well.	Attribution	OM	He felt well?	Yes/No Question
		YM	One cannot doubt it.	Explain
The metals are not suddenly deposited in the ores. It is the patient work of countless ages.	Contrast	OM	Are the metals suddenly deposited in the ores?	Yes/No Question
		YM	No --	Answer No
		YM	it is the patient work of countless ages.	Explain

Reactive dialogue systems

Beyond ELIZA - Harvesting generation rules from text

CODA (Piwek & Stoyanchev, 2010; Kuyten et al., 2012)

1. Create a parallel monologue – dialogue corpus
2. Annotate the corpus with discourse relations and dialogue acts
3. Extract monologue-to-dialogue generation patterns
4. Apply the patterns to a new text (monologue), obtaining a set of generation rules
5. Use the generation rules in a dialogue system

   ... or for generating scripted dialogues from text

See also Yao et al. (2012)

Agenda-driven systems

- Fixed-task agenda
- Selected-task agenda
- Joint-task agenda

Agenda-driven systems – fixed task

Philips train timetable information system (Aust et al., 1995)

1	System	Good morning, this is the automatic time-table information service. How can I help you?
2	Caller	Good morning. I need a connection from Hamburg to Munich.
3	System	When would you like to go from Hamburg to Munich?
4	Caller	Next Friday, er, at eight o'clock.
5	System	So you want to travel on Sunday the 25th of September at 8am?
6	Caller	No, on Friday.
7	System	So you want to travel on Friday the 23rd of September at 8am?
8	Caller	Exactly.
9	System	There is the following connection: With Intercity Express 793, departure from Hamburg at 8:02, arrival in Munich at 14:05. Would you like me to repeat the connection?
10	Caller	No, but I need a return trip on Sunday.
11	System	(…)

Bobrow et al. (1977) GUS, A Frame-driven Dialogue System

Agenda-driven systems – fixed task

The store

Private agenda

1. ask for the place of departure
2. ask for the destination
3. ask for the time of travel
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure		
Destination		
Travel time		

Agenda-driven systems – fixed task

The store

Private agenda

1. ask for the place of departure
2. ask for the destination
3. ask for the time of travel
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure		
Destination		
Travel time		

System Good morning, this is the automatic time-table information service. How can I help you?

Agenda-driven systems – fixed task

The store

Private agenda

1. ask for the place of departure
2. ask for the destination
3. ask for the time of travel
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure		
Destination		
Travel time		

Caller Good morning. I need a connection from Hamburg to Munich.

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ask for the time of travel
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Unconf.
Destination	Munich	Unconf.
Travel time		

Caller Good morning. I need a connection from Hamburg to Munich.

Agenda-driven systems – fixed task

The generation rule

Retrieve the next item on the agenda
retrieve any unconfirmed slots

then

formulate an utterance that addresses the next agenda item, if possible, whilst confirming any unconfirmed slots

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ask for the time of travel
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Unconf.
Destination	Munich	Unconf.
Travel time		

System When would you like to go from Hamburg to Munich?

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ~~ask for the time of travel~~
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 25, 8am, Sunday	Unconf.

Caller Next Friday, er, at eight o’clock.

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ~~ask for the time of travel~~
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 25, 8am, Sunday	Unconf.

System   So you want to travel on Sunday the 25th of September at 8am?

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ~~ask for the time of travel~~
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 23, 8am, Friday	Unconf.

Caller No, on Friday.

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ~~ask for the time of travel~~
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 23, 8am, Friday	Unconf.

System So you want to travel on Friday the 23rd of September at 8am?

Agenda-driven systems – fixed task

The store

Private agenda

1. ~~ask for the place of departure~~
2. ~~ask for the destination~~
3. ~~ask for the time of travel~~
4. provide the connection

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 23, 8am, Friday	Confirmed

Caller Exactly.

Agenda-driven systems – fixed task

The generation rule

Retrieve the next item on the agenda
retrieve any unconfirmed slots

then

formulate an utterance that addresses the next agenda item, if possible, whilst confirming any unconfirmed slots

Agenda-driven systems – fixed task

The store

Private agenda

- ~~ask for the place of departure~~
- ~~ask for the destination~~
- ~~ask for the time of travel~~
- ~~provide the connection~~

Common

Attribute	Value	Confidence
Point of departure	Hamburg	Confirmed
Destination	Munich	Confirmed
Travel time	Sept 23, 8am, Friday	Confirmed

System There is the following connection: With Intercity Express 793, departure from Hamburg at 8:02, arrival in Munich at 14:05.

Intermezzo 1 - Evaluation

PARADISE: PARAdigm for DIalogue System Evaluation (Walker, Litman, Kamm & Abella, 1997: 272)

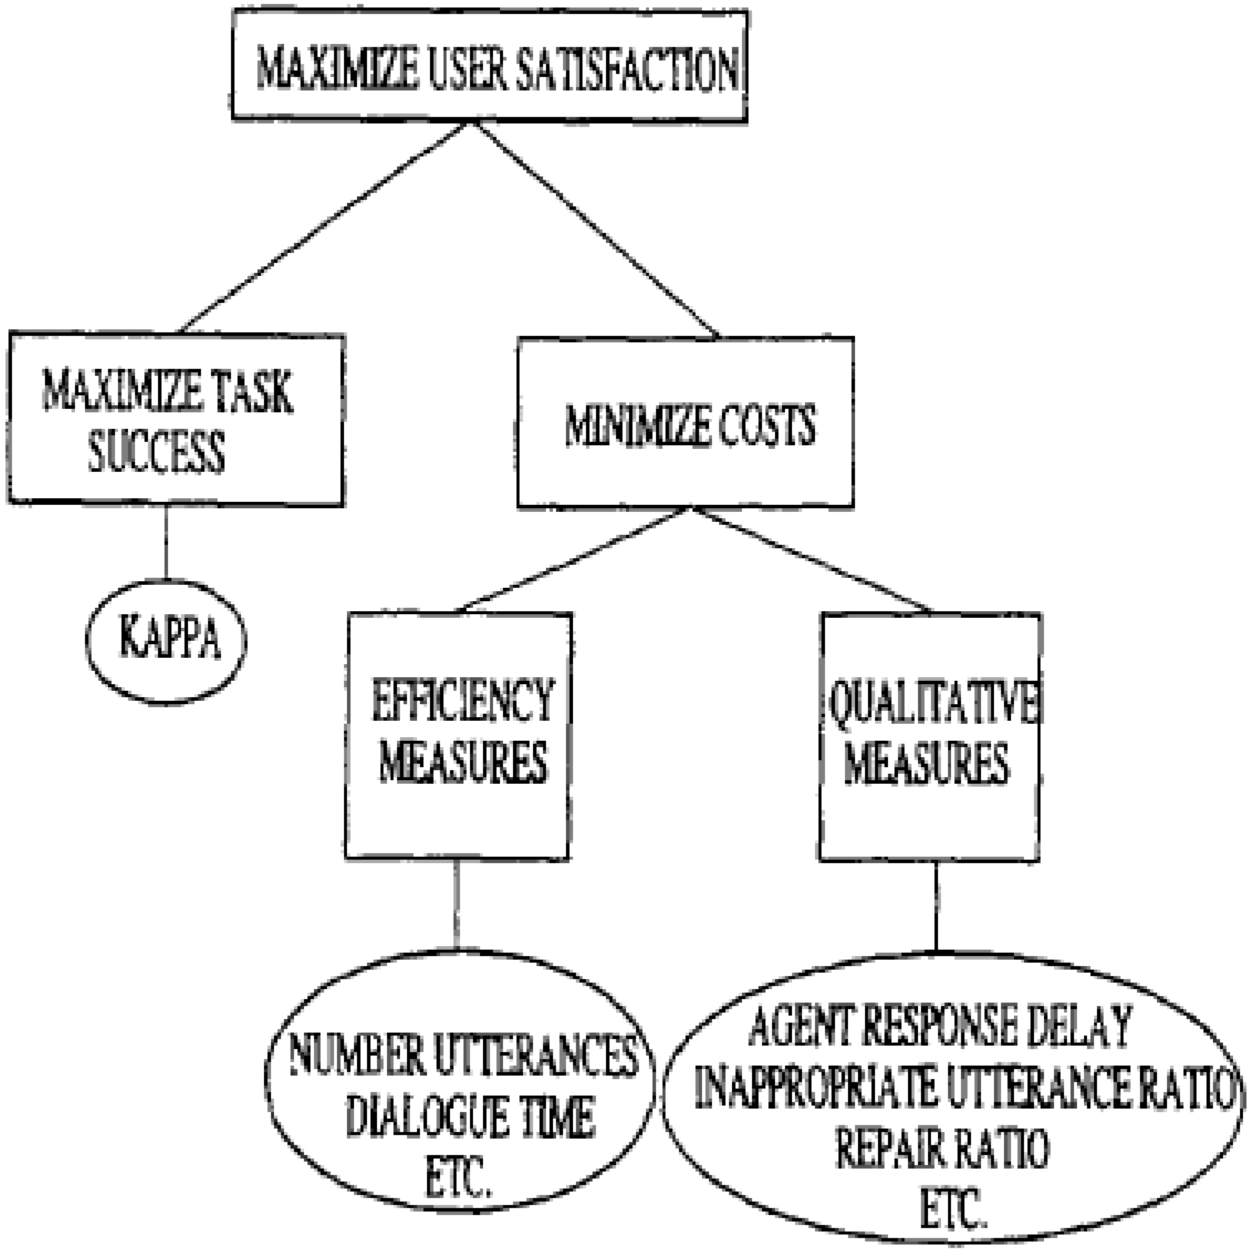

Intermezzo 1 - Evaluation

PARADISE: PARAdigm for DIalogue System Evaluation
(Walker, Litman, Kamm & Abella, 1997: 273) – Task success

attribute	possible values	information flow
depart-city (DC)	Milano, Roma, Torino, Trento	to agent
arrival-city (AC)	Milano, Roma, Torino, Trento	to agent
depart-range (DR)	morning,evening	to agent
depart-time (DT)	6am,8am,6pm,8pm	to user

attribute	actual value
depart-city	Torino
arrival-city	Milano
depart-range	evening
depart-time	8pm

Intermezzo 1 - Evaluation

PARADISE: PARAdigm for DIalogue System Evaluation (Walker, Litman, Kamm & Abella, 1997: 273) – Task success

- Naive approach: calculate agreement between actual slot fillers and scenario in question. So, we could say, for instance, “80% of the dialogues were successful”

Intermezzo 1 - Evaluation

PARADISE: PARAdigm for DIalogue System Evaluation
(Walker, Litman, Kamm & Abella, 1997: 273) – Task success

	KEY													
	DEPART-CITY				ARRIVAL-CITY				DEPART-RANGE		DEPART-TIME			
DATA	v1	v2	v3	v4	v5	v6	v7	v8	v9	v10	v11	v12	v13	v14
v1	**22**		1		3									
v2		**29**												
v3	4		**16**	4			1							
v4	1	1	5	**11**			1							
v5	3				**20**									
v6						**22**								
v7			2		1	1	**20**	5						
v8			1		1	2	8	**15**						
v9									**45**	10				
v10									5	**40**				
v11											**20**		2	
v12											1	**19**	2	4
v13											2		**18**	
v14											2	6	3	**21**
sum	30	30	25	15	25	25	30	20	50	50	25	25	25	25

$$\kappa = \frac{P(A) - P(E)}{1 - P(E)}$$

Intermezzo 1 - Evaluation

PARADISE: PARAdigm for DIalogue System Evaluation (Walker, Litman, Kamm & Abella, 1997: 273) – Overall

Performance = (w_S * Task success) – (w_C * Normalised Cost)

Find values for weights w_S and w_C via multiple linear regression.

This now allows us to calculate overall performance/predict user satisfaction from Task success and Cost.

Intermezzo 2 - Learning

Dialogue manager

Intermezzo 2 - Learning

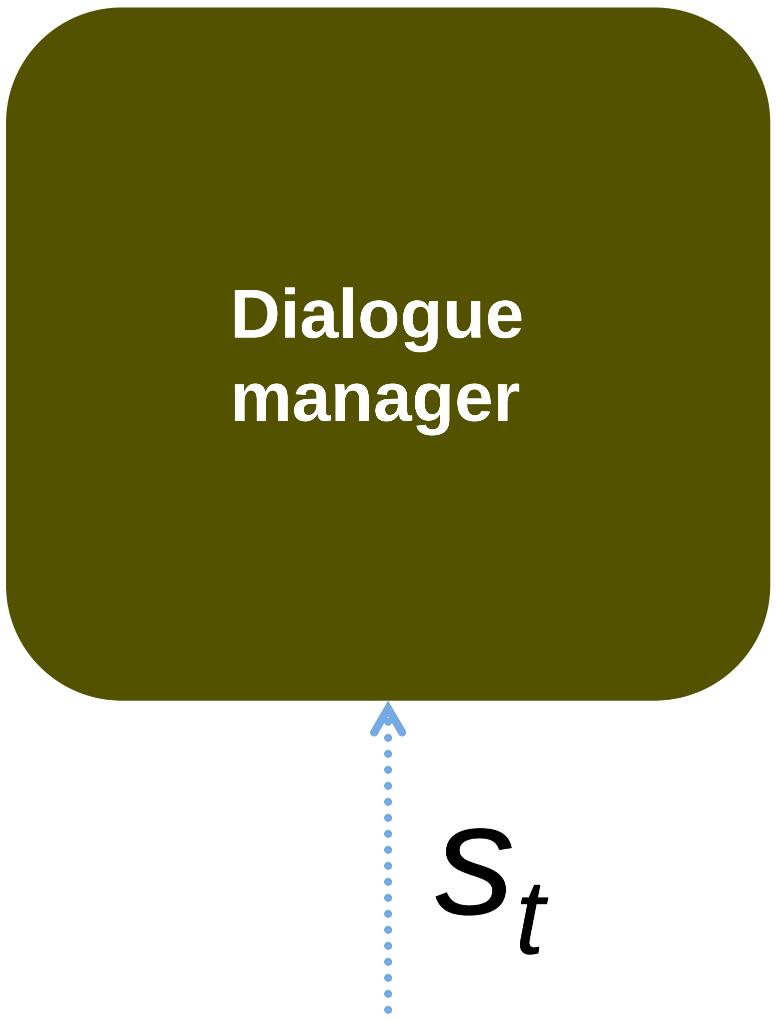

Intermezzo 2 - Learning

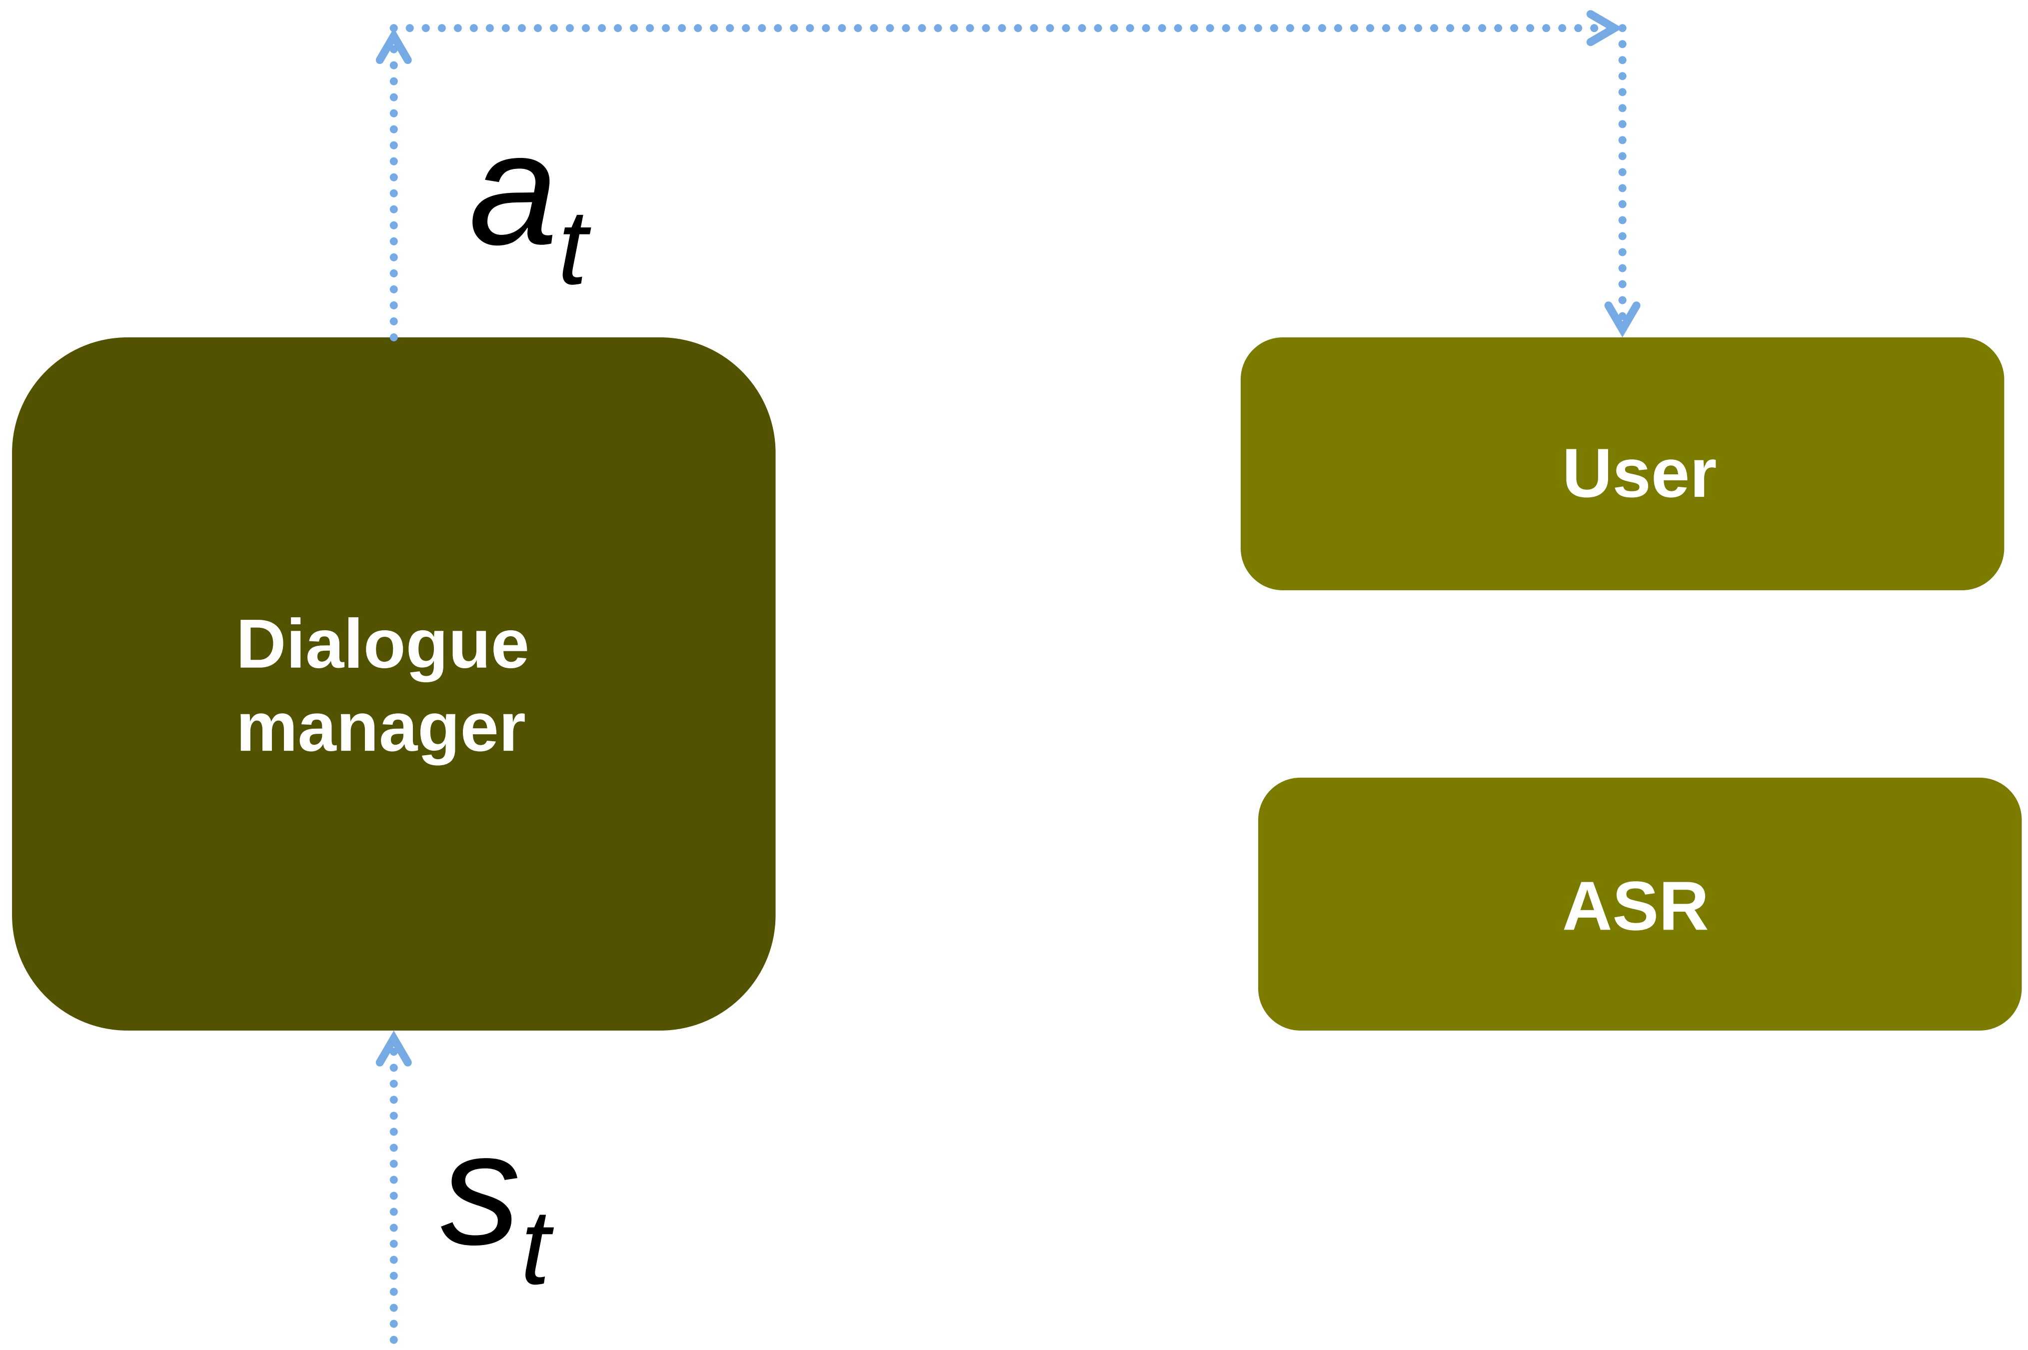

Intermezzo 2 - Learning

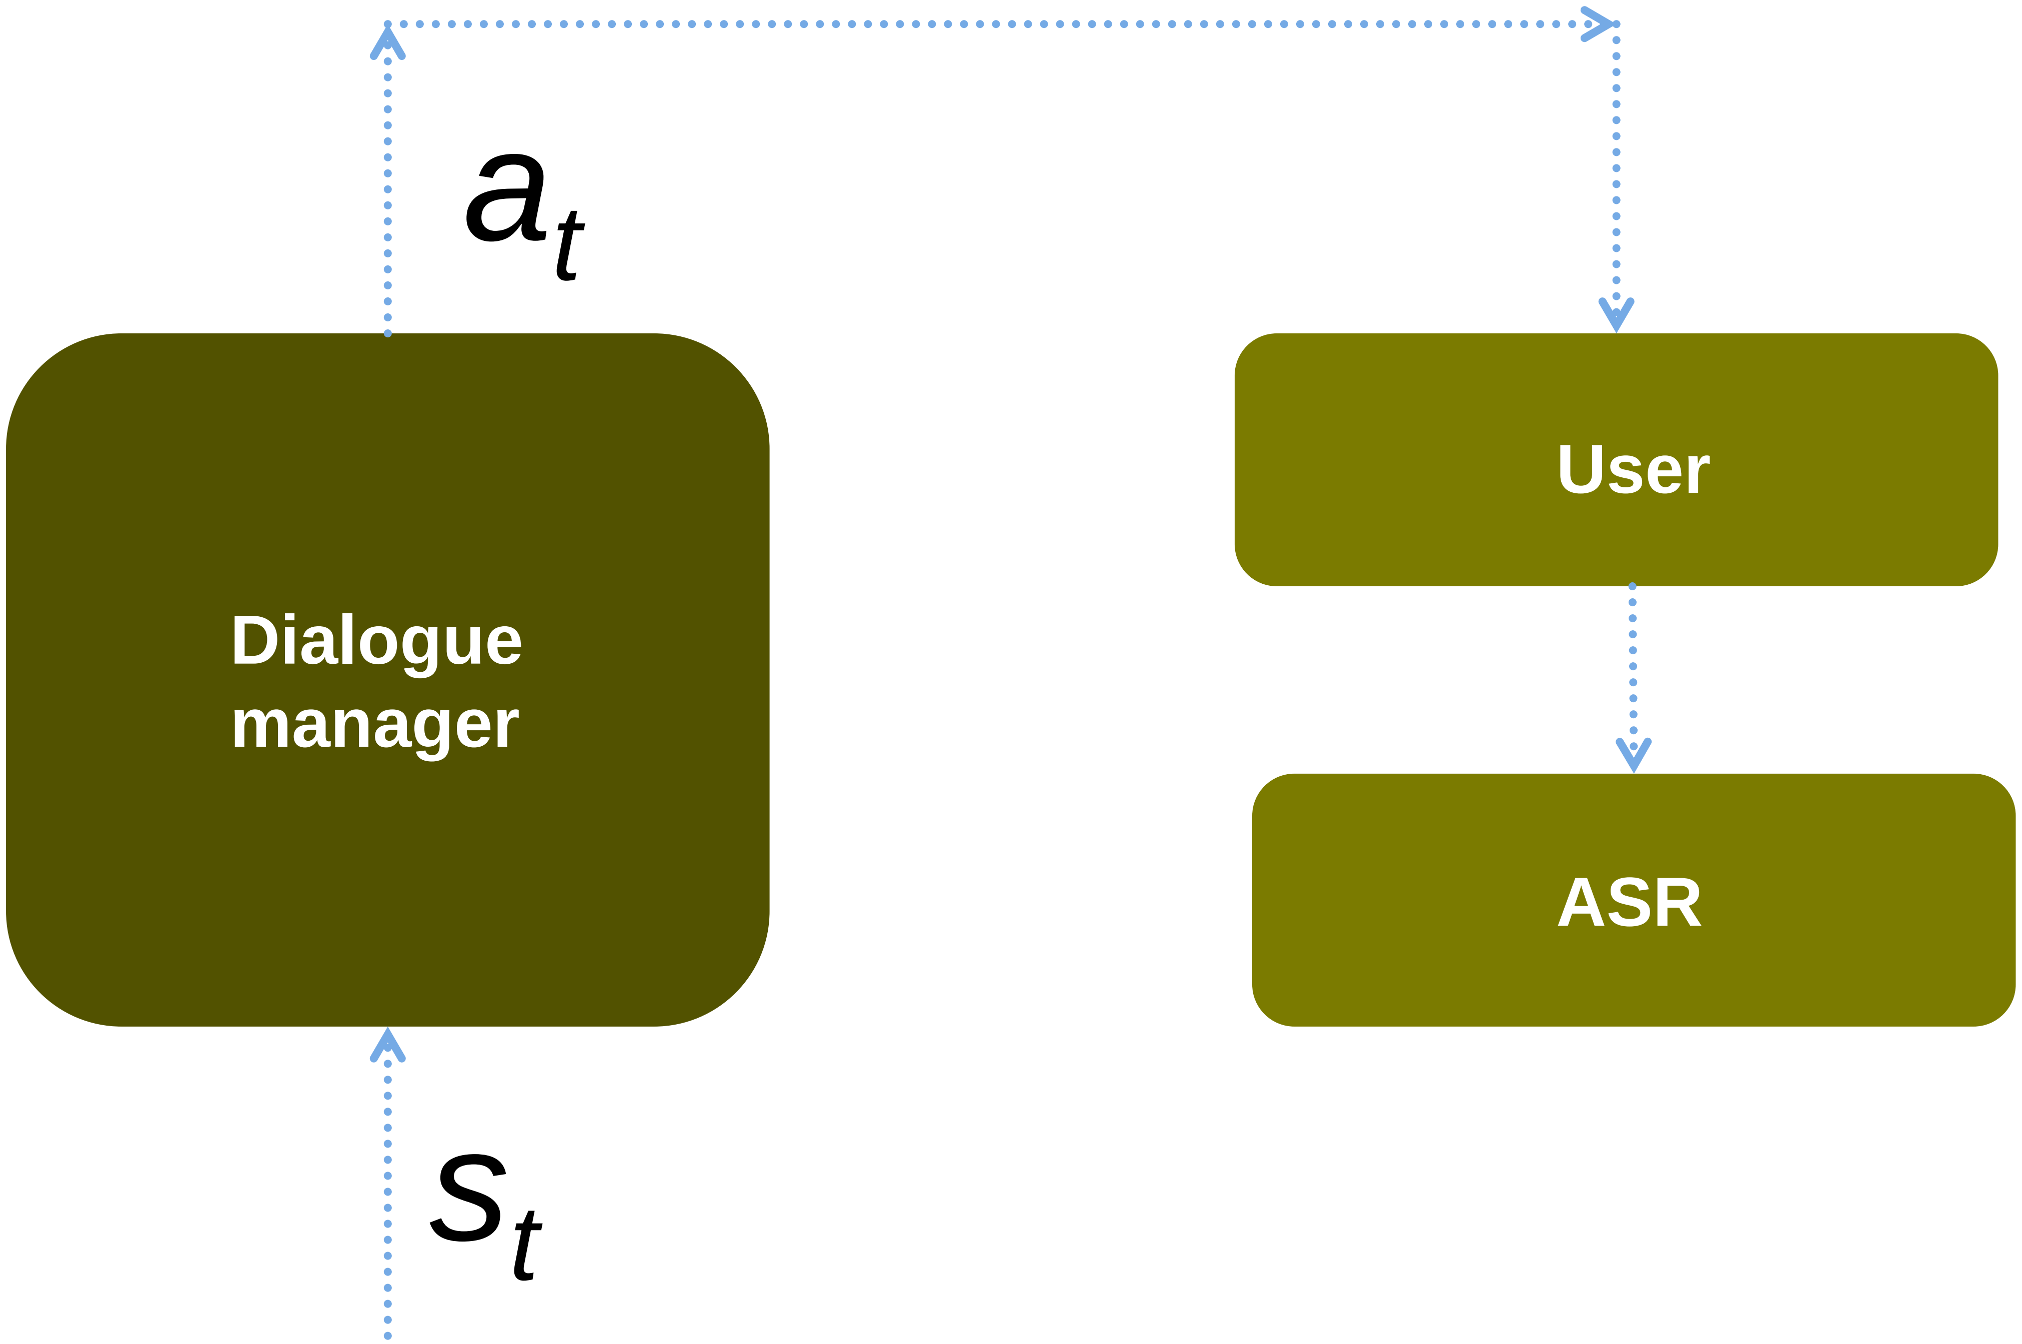

Intermezzo 2 - Learning

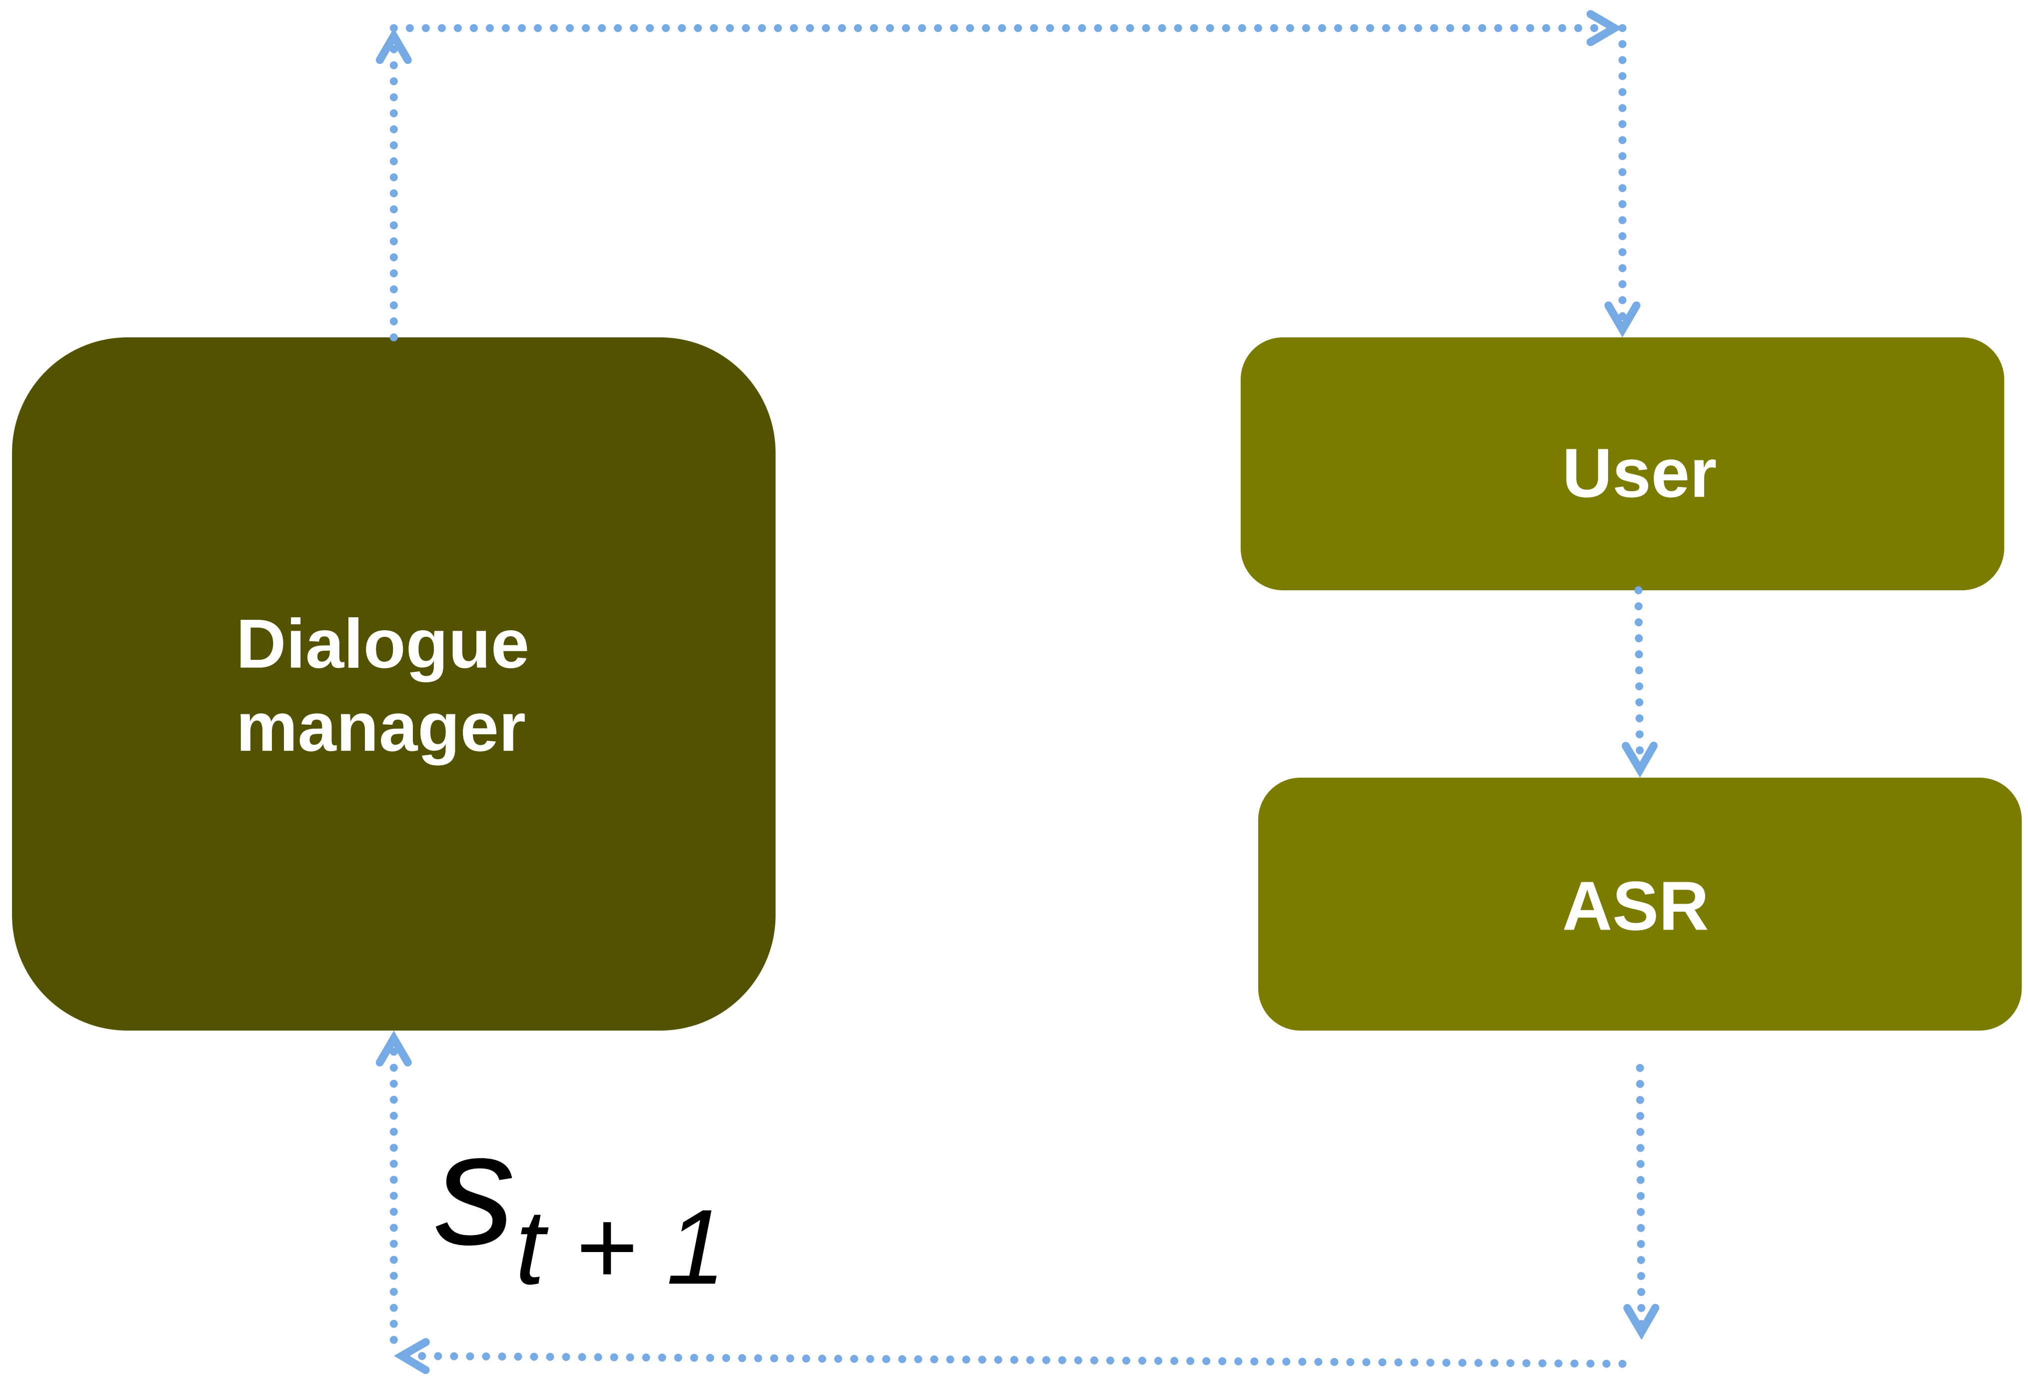

Intermezzo 2 - Learning

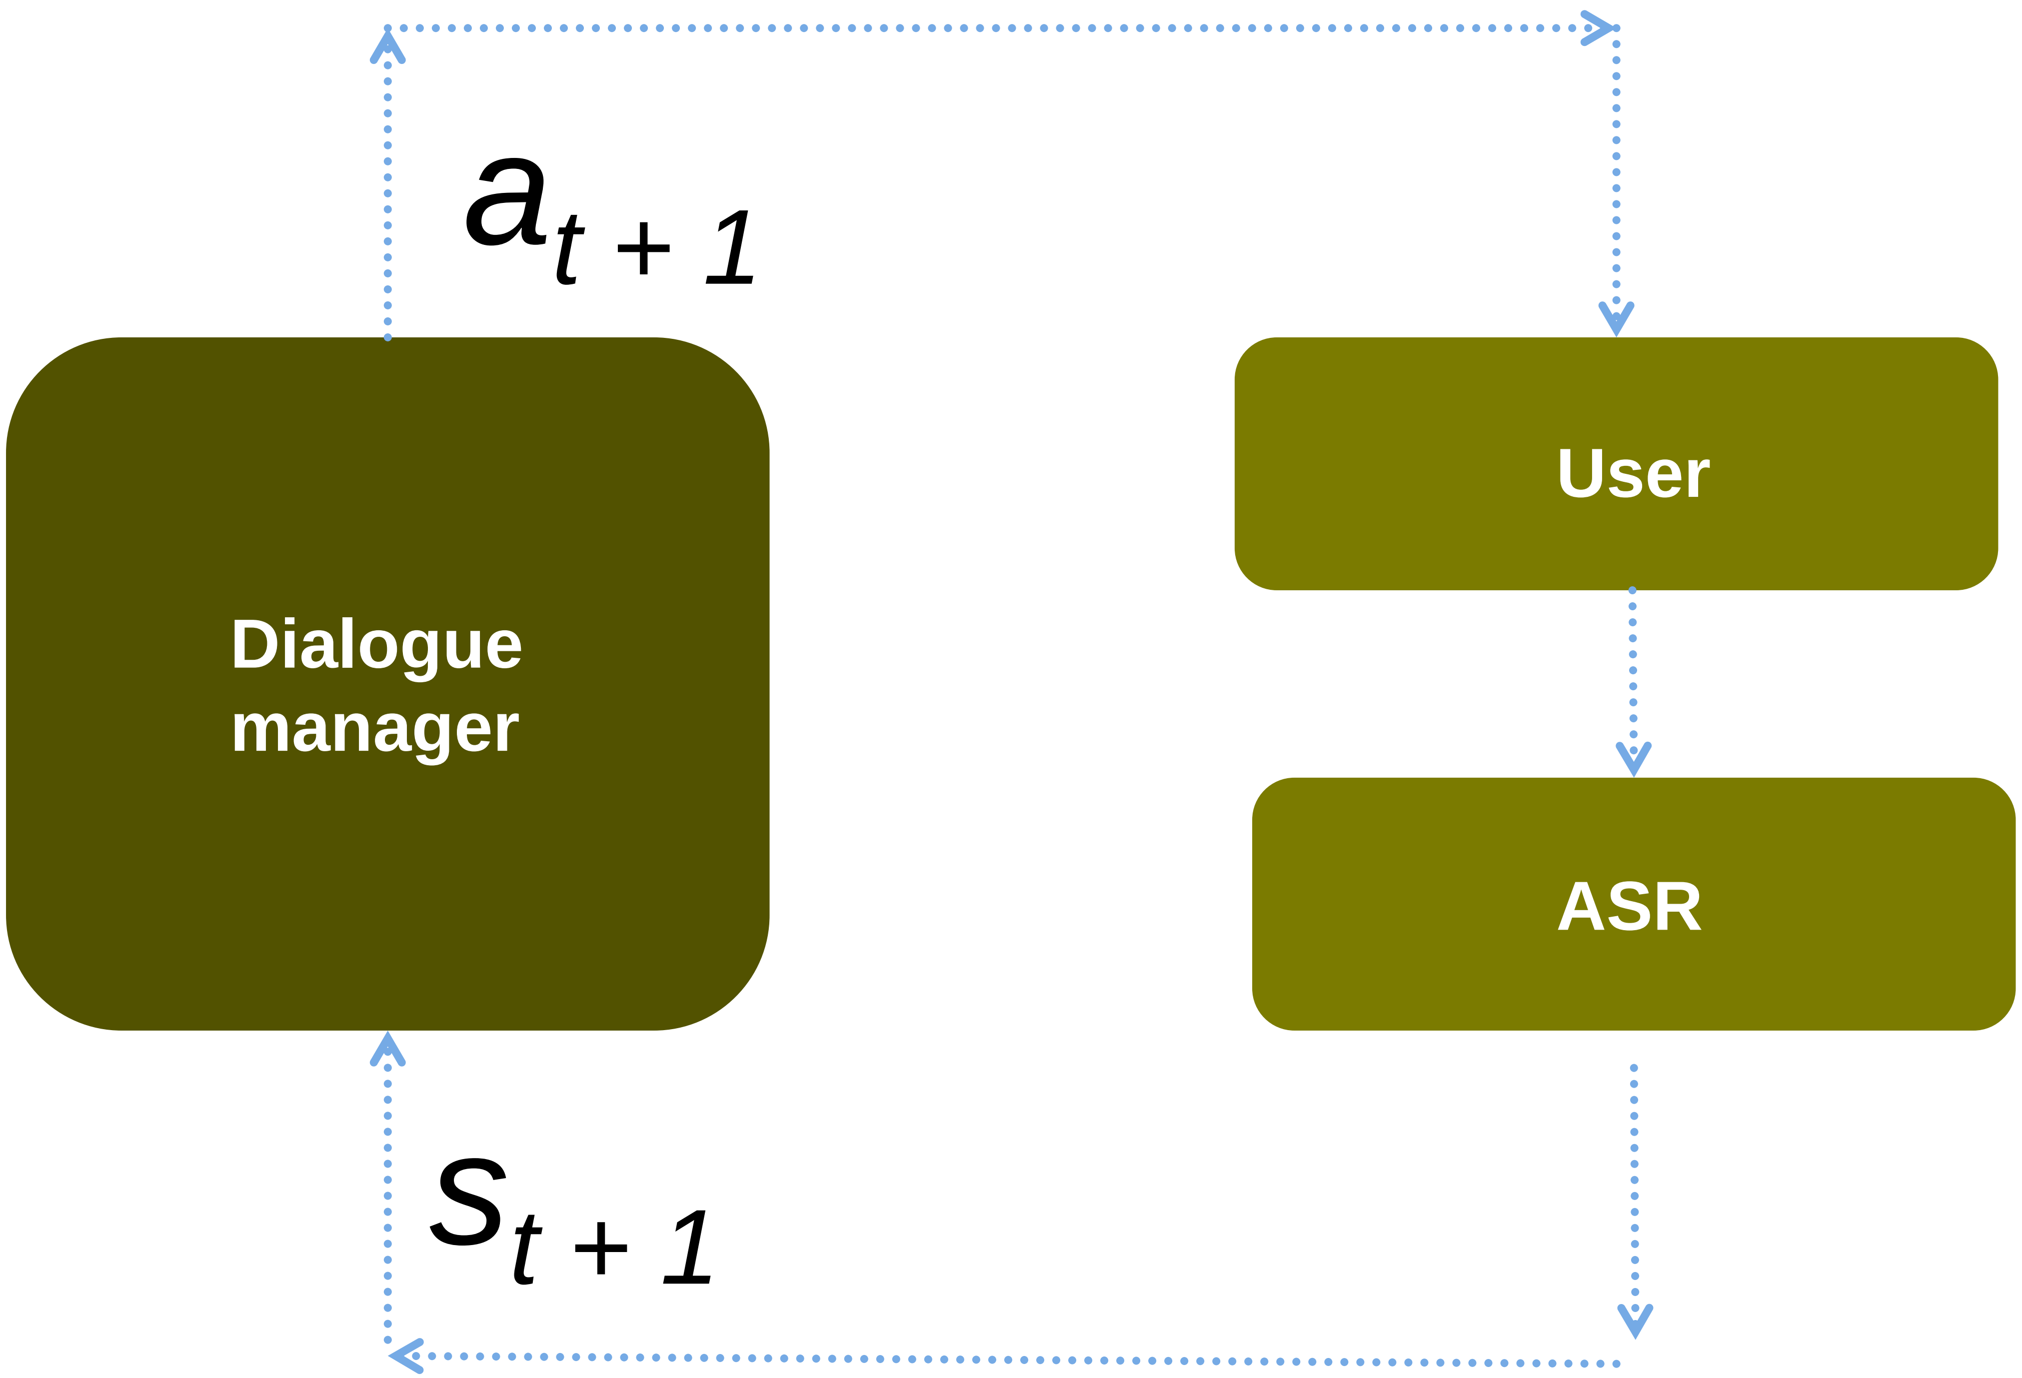

Intermezzo 2 - Learning

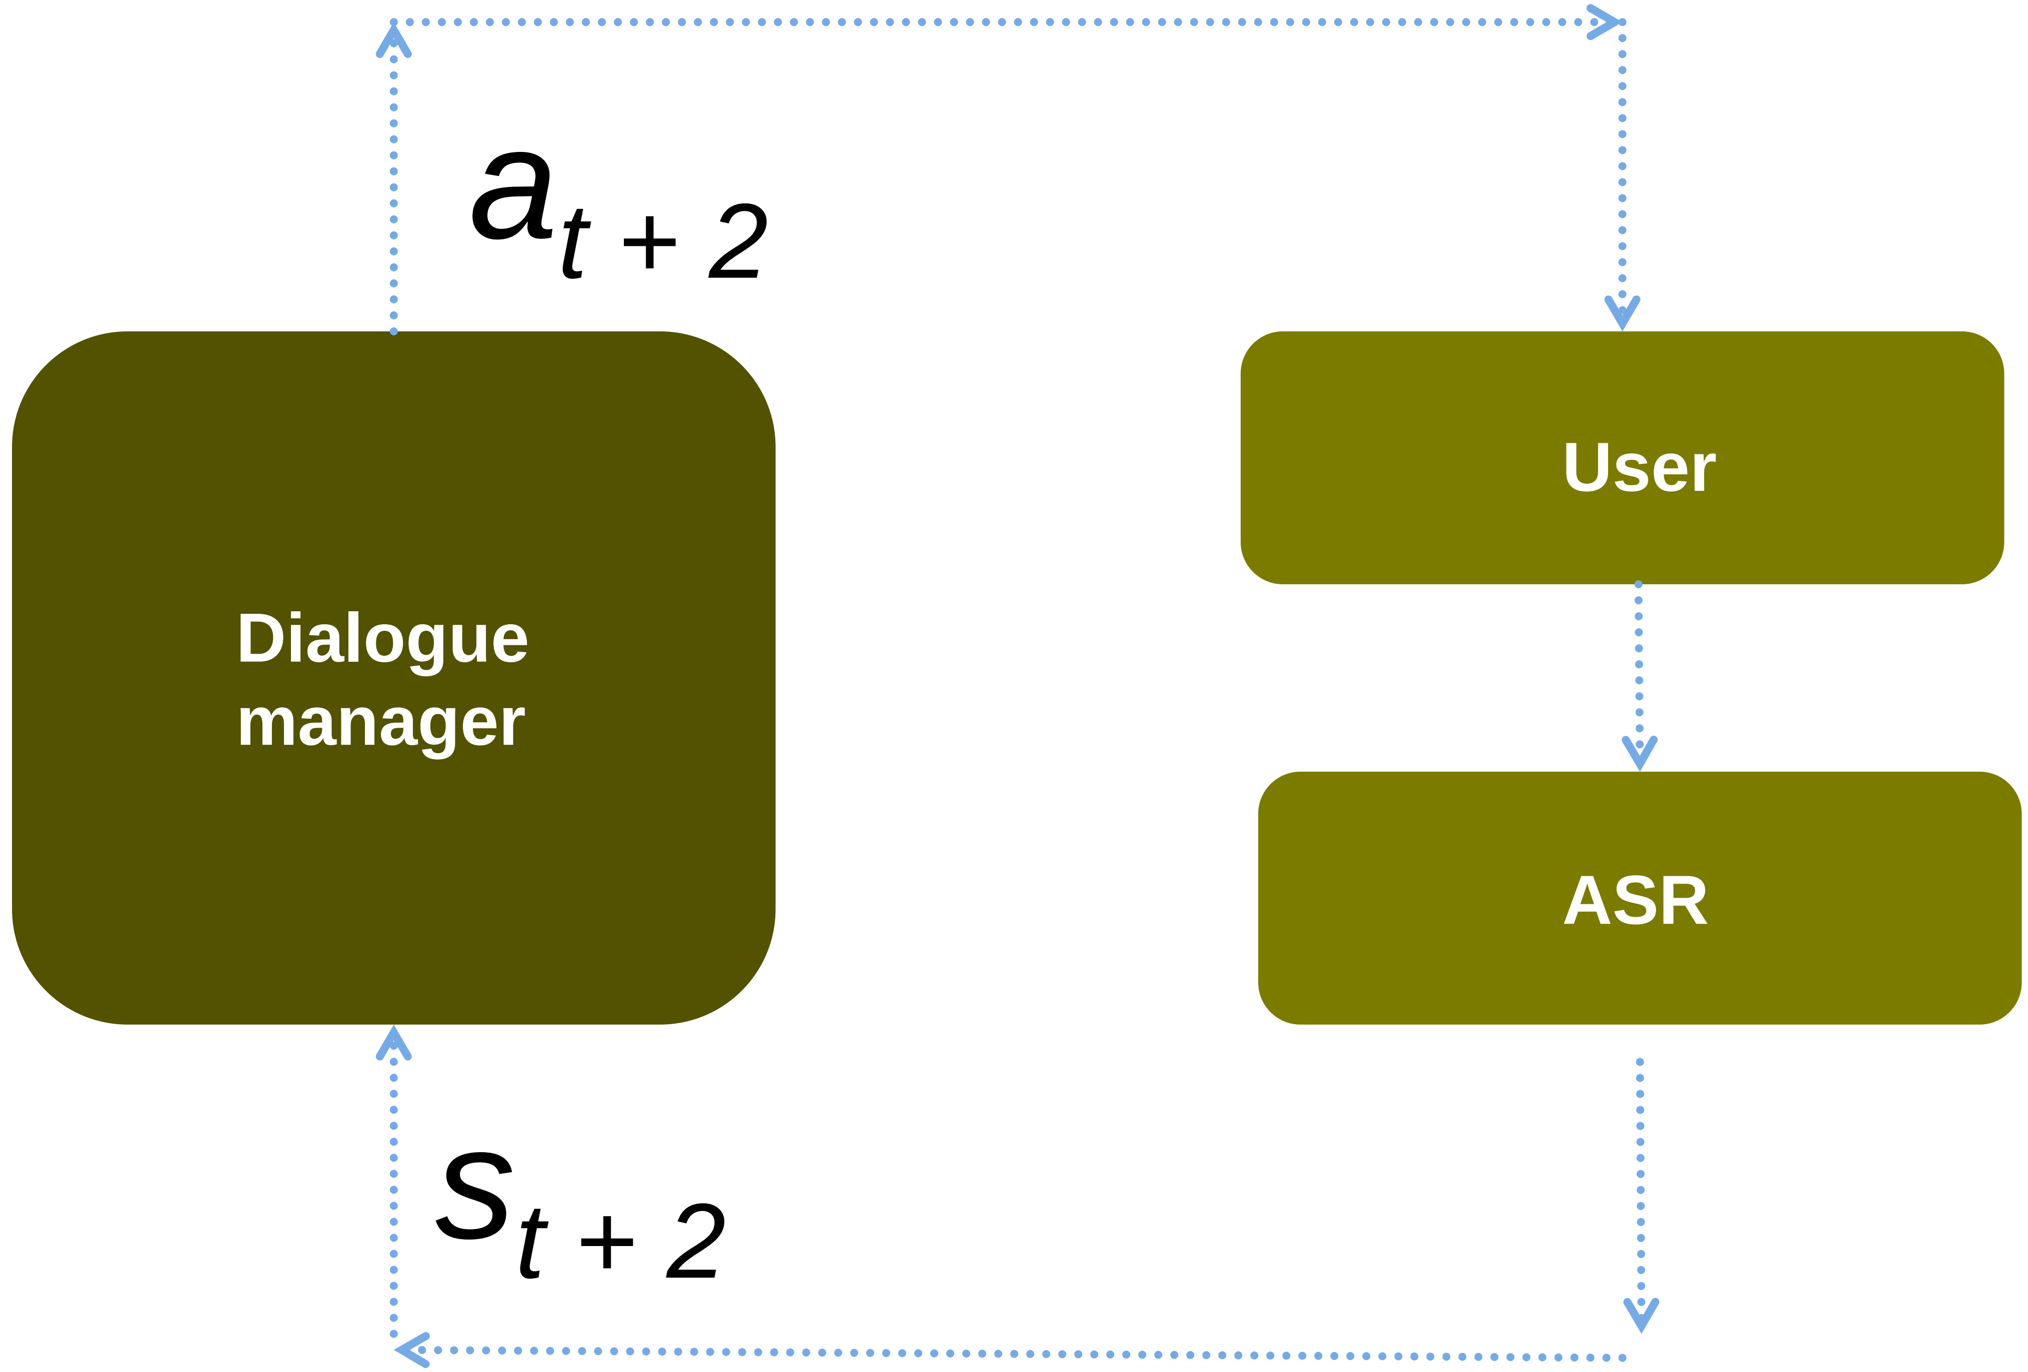

Intermezzo 2 - Learning

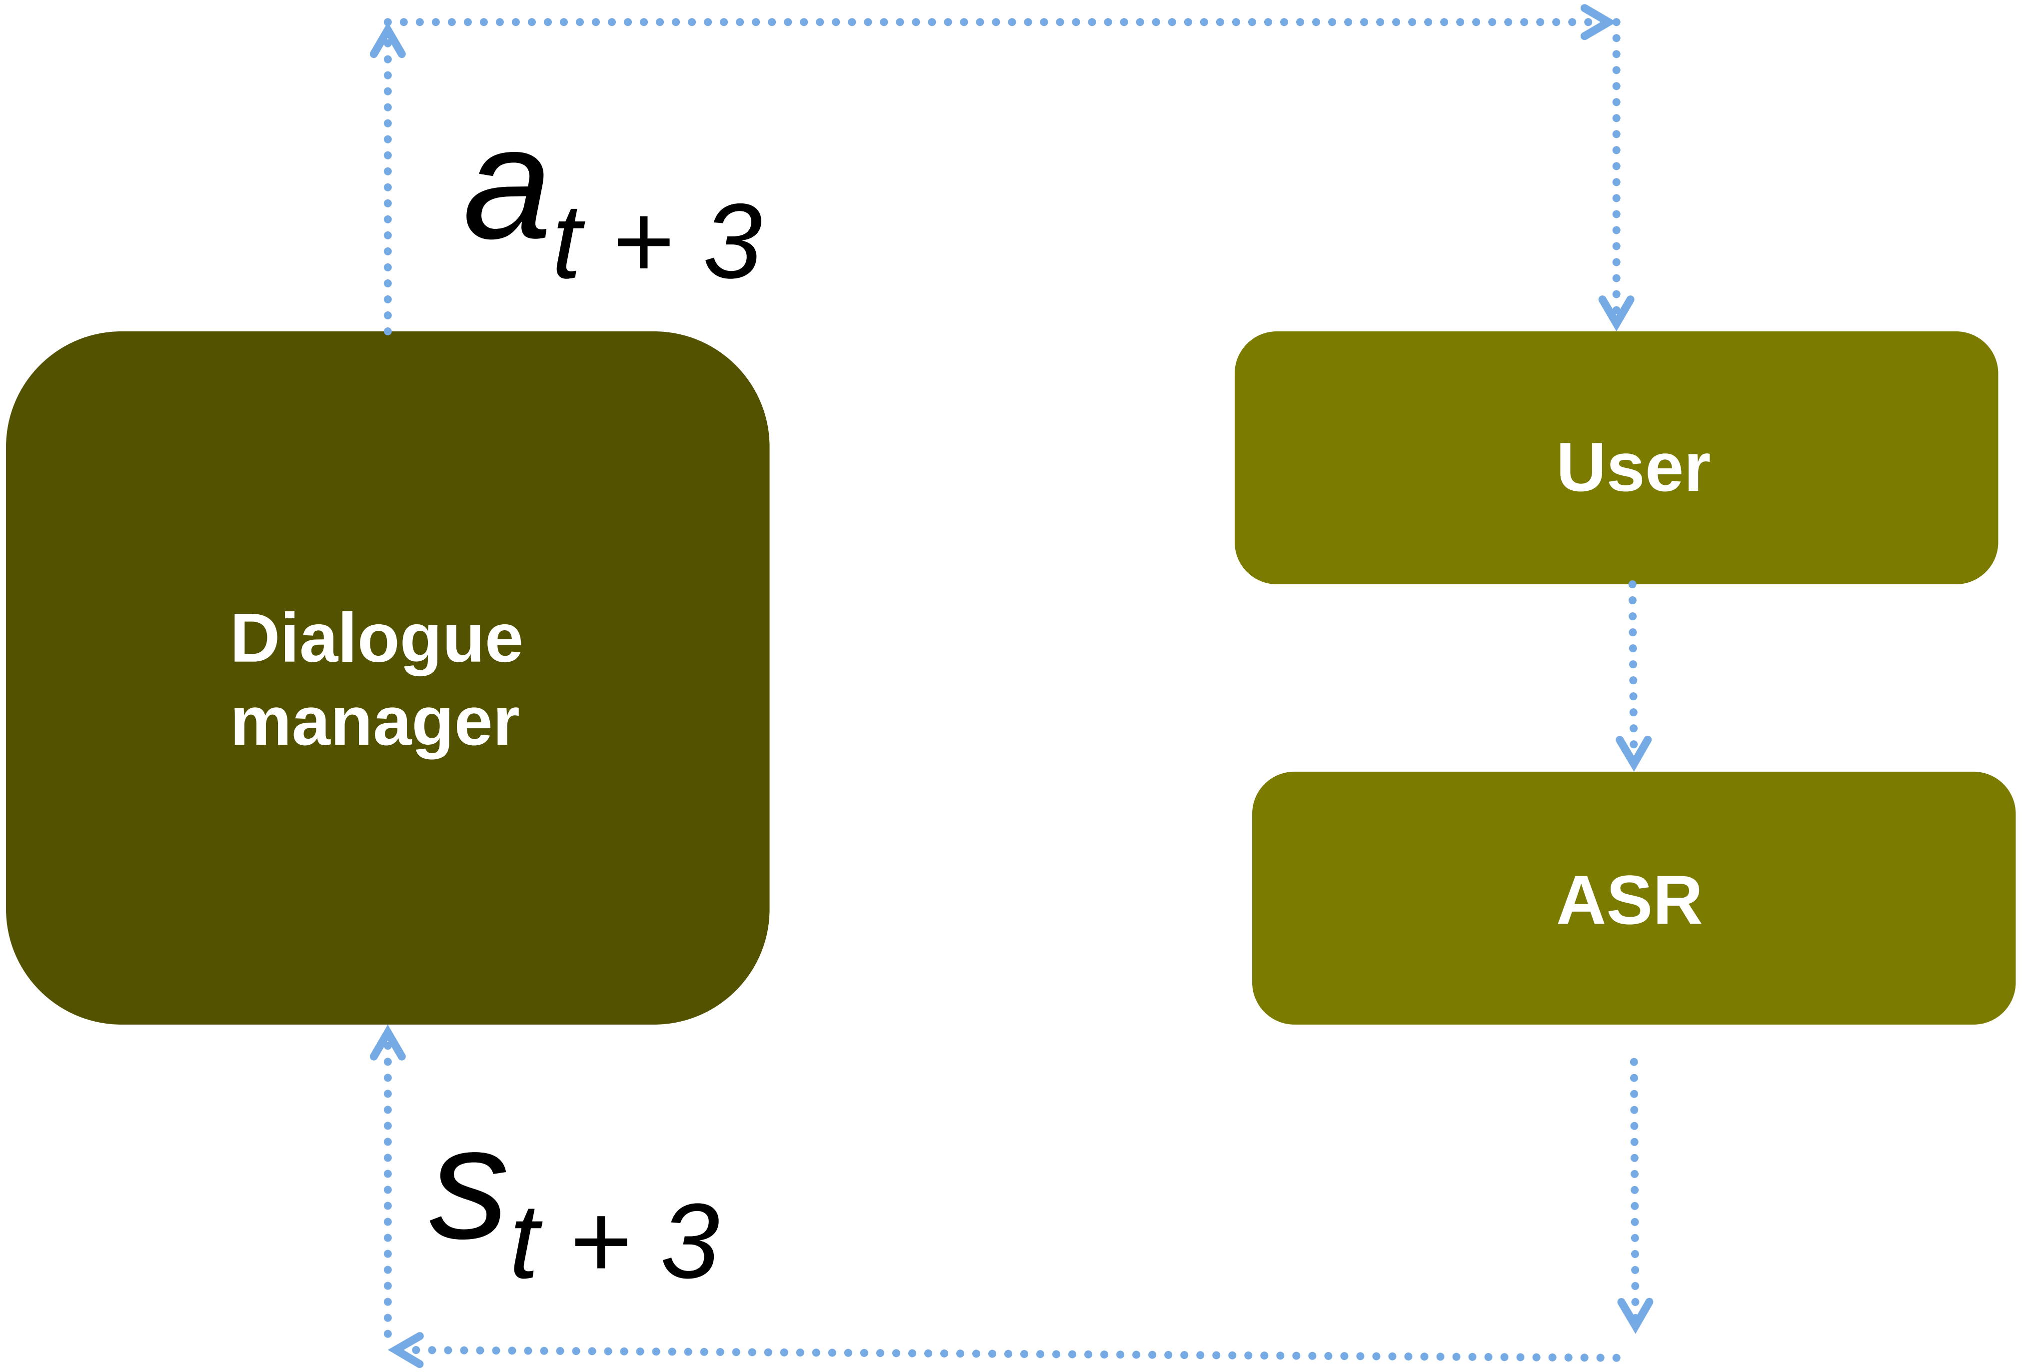

Intermezzo 2 - Learning

Markov decision process

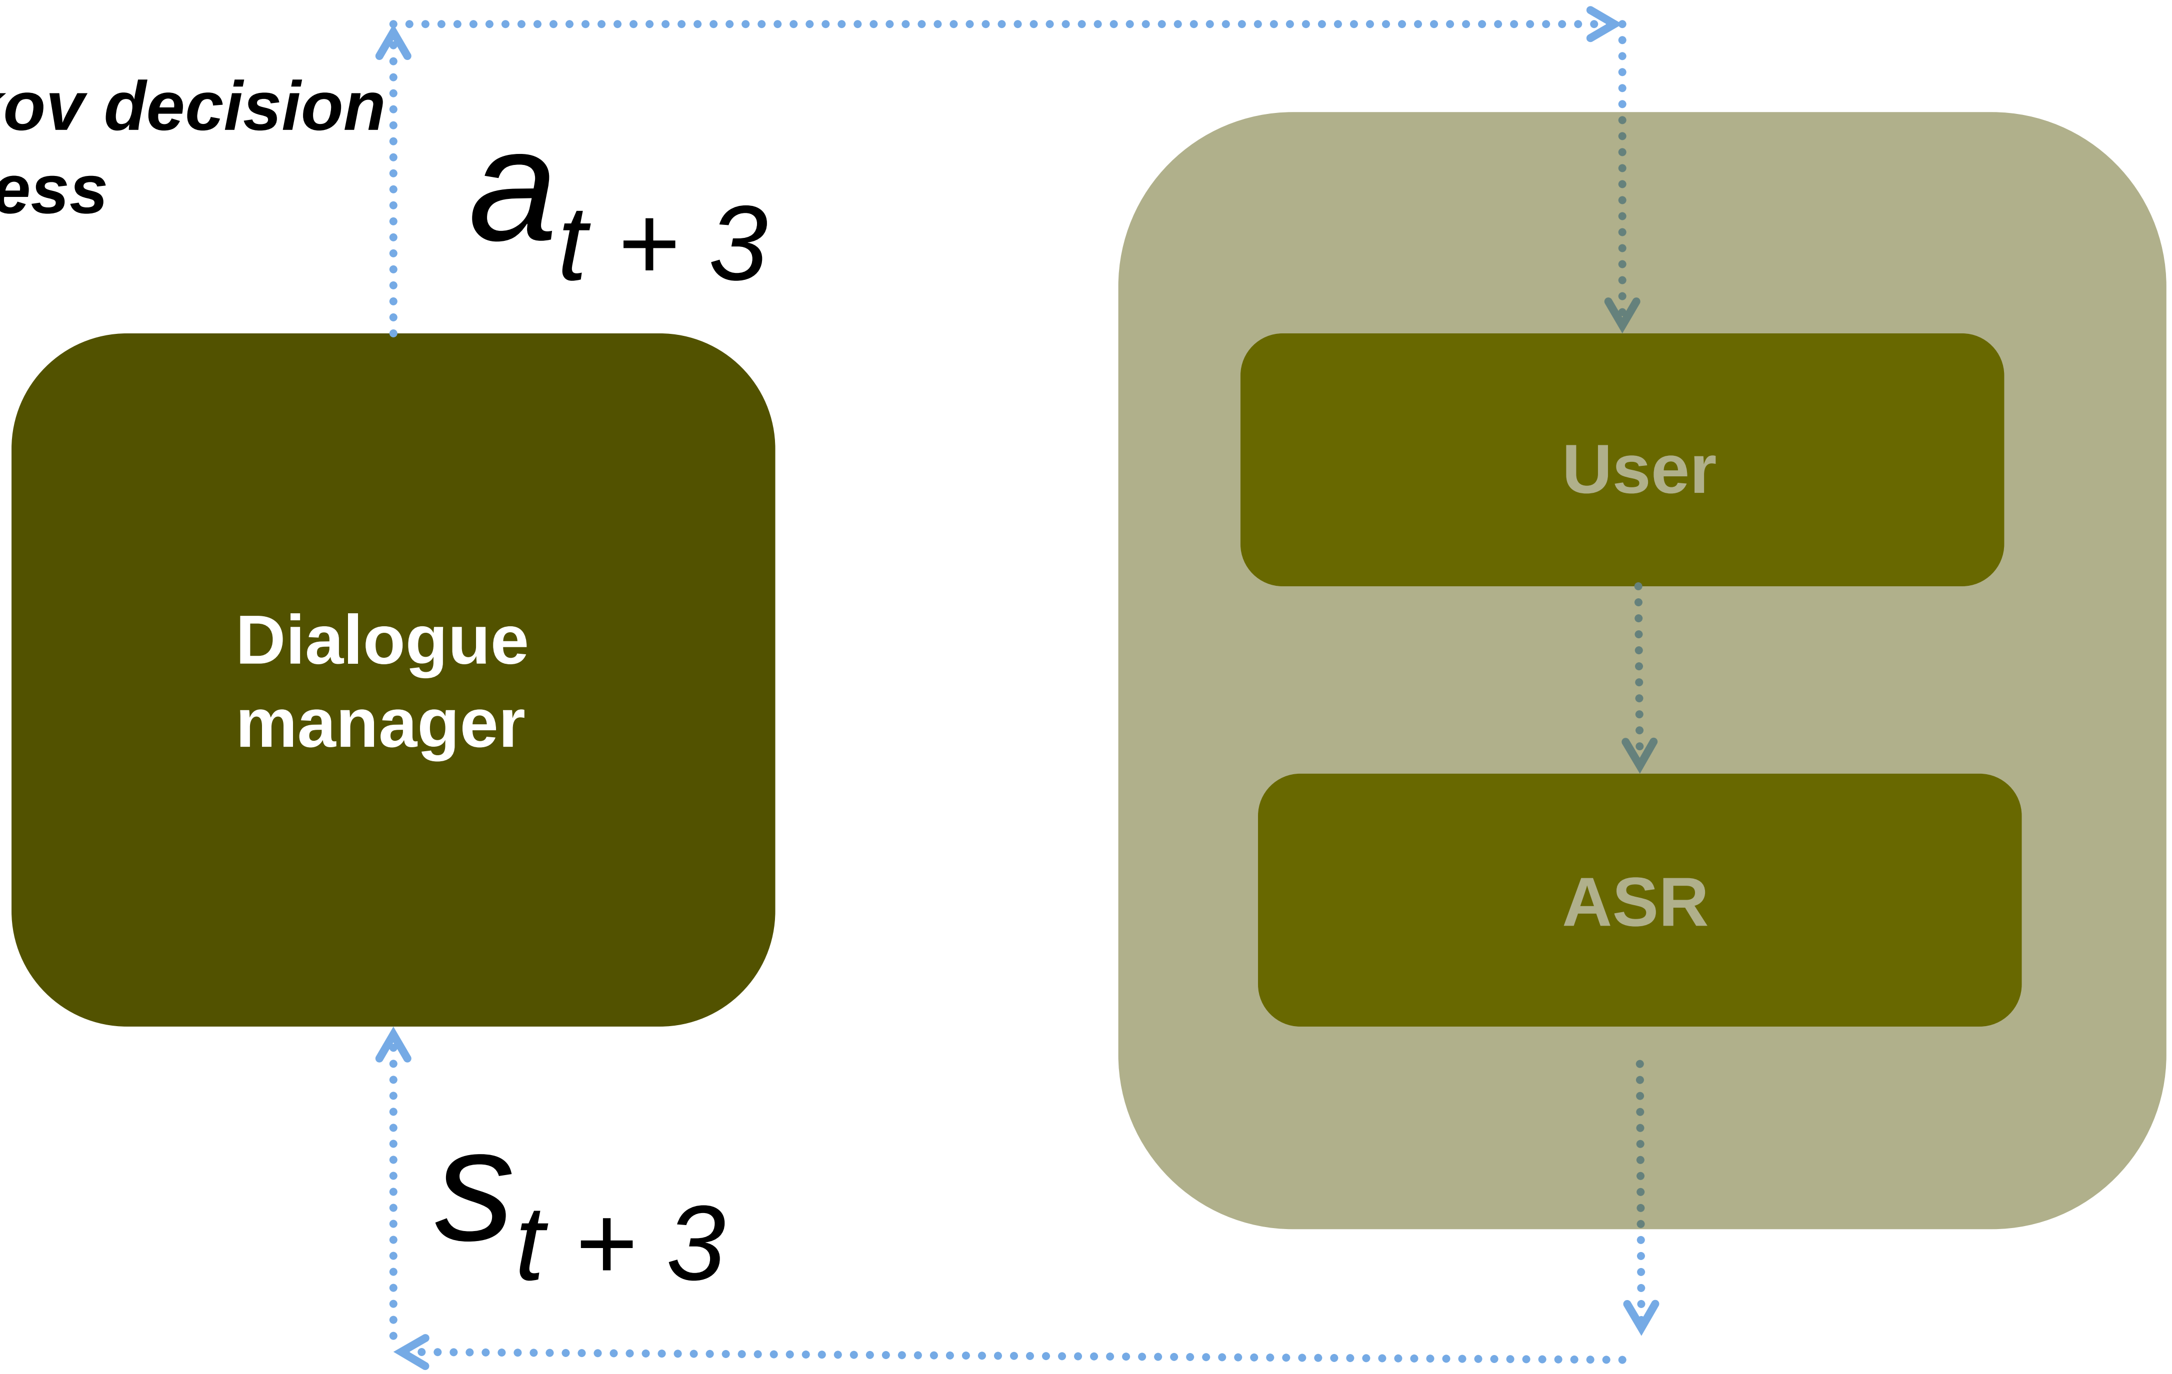

Intermezzo 2 - Learning

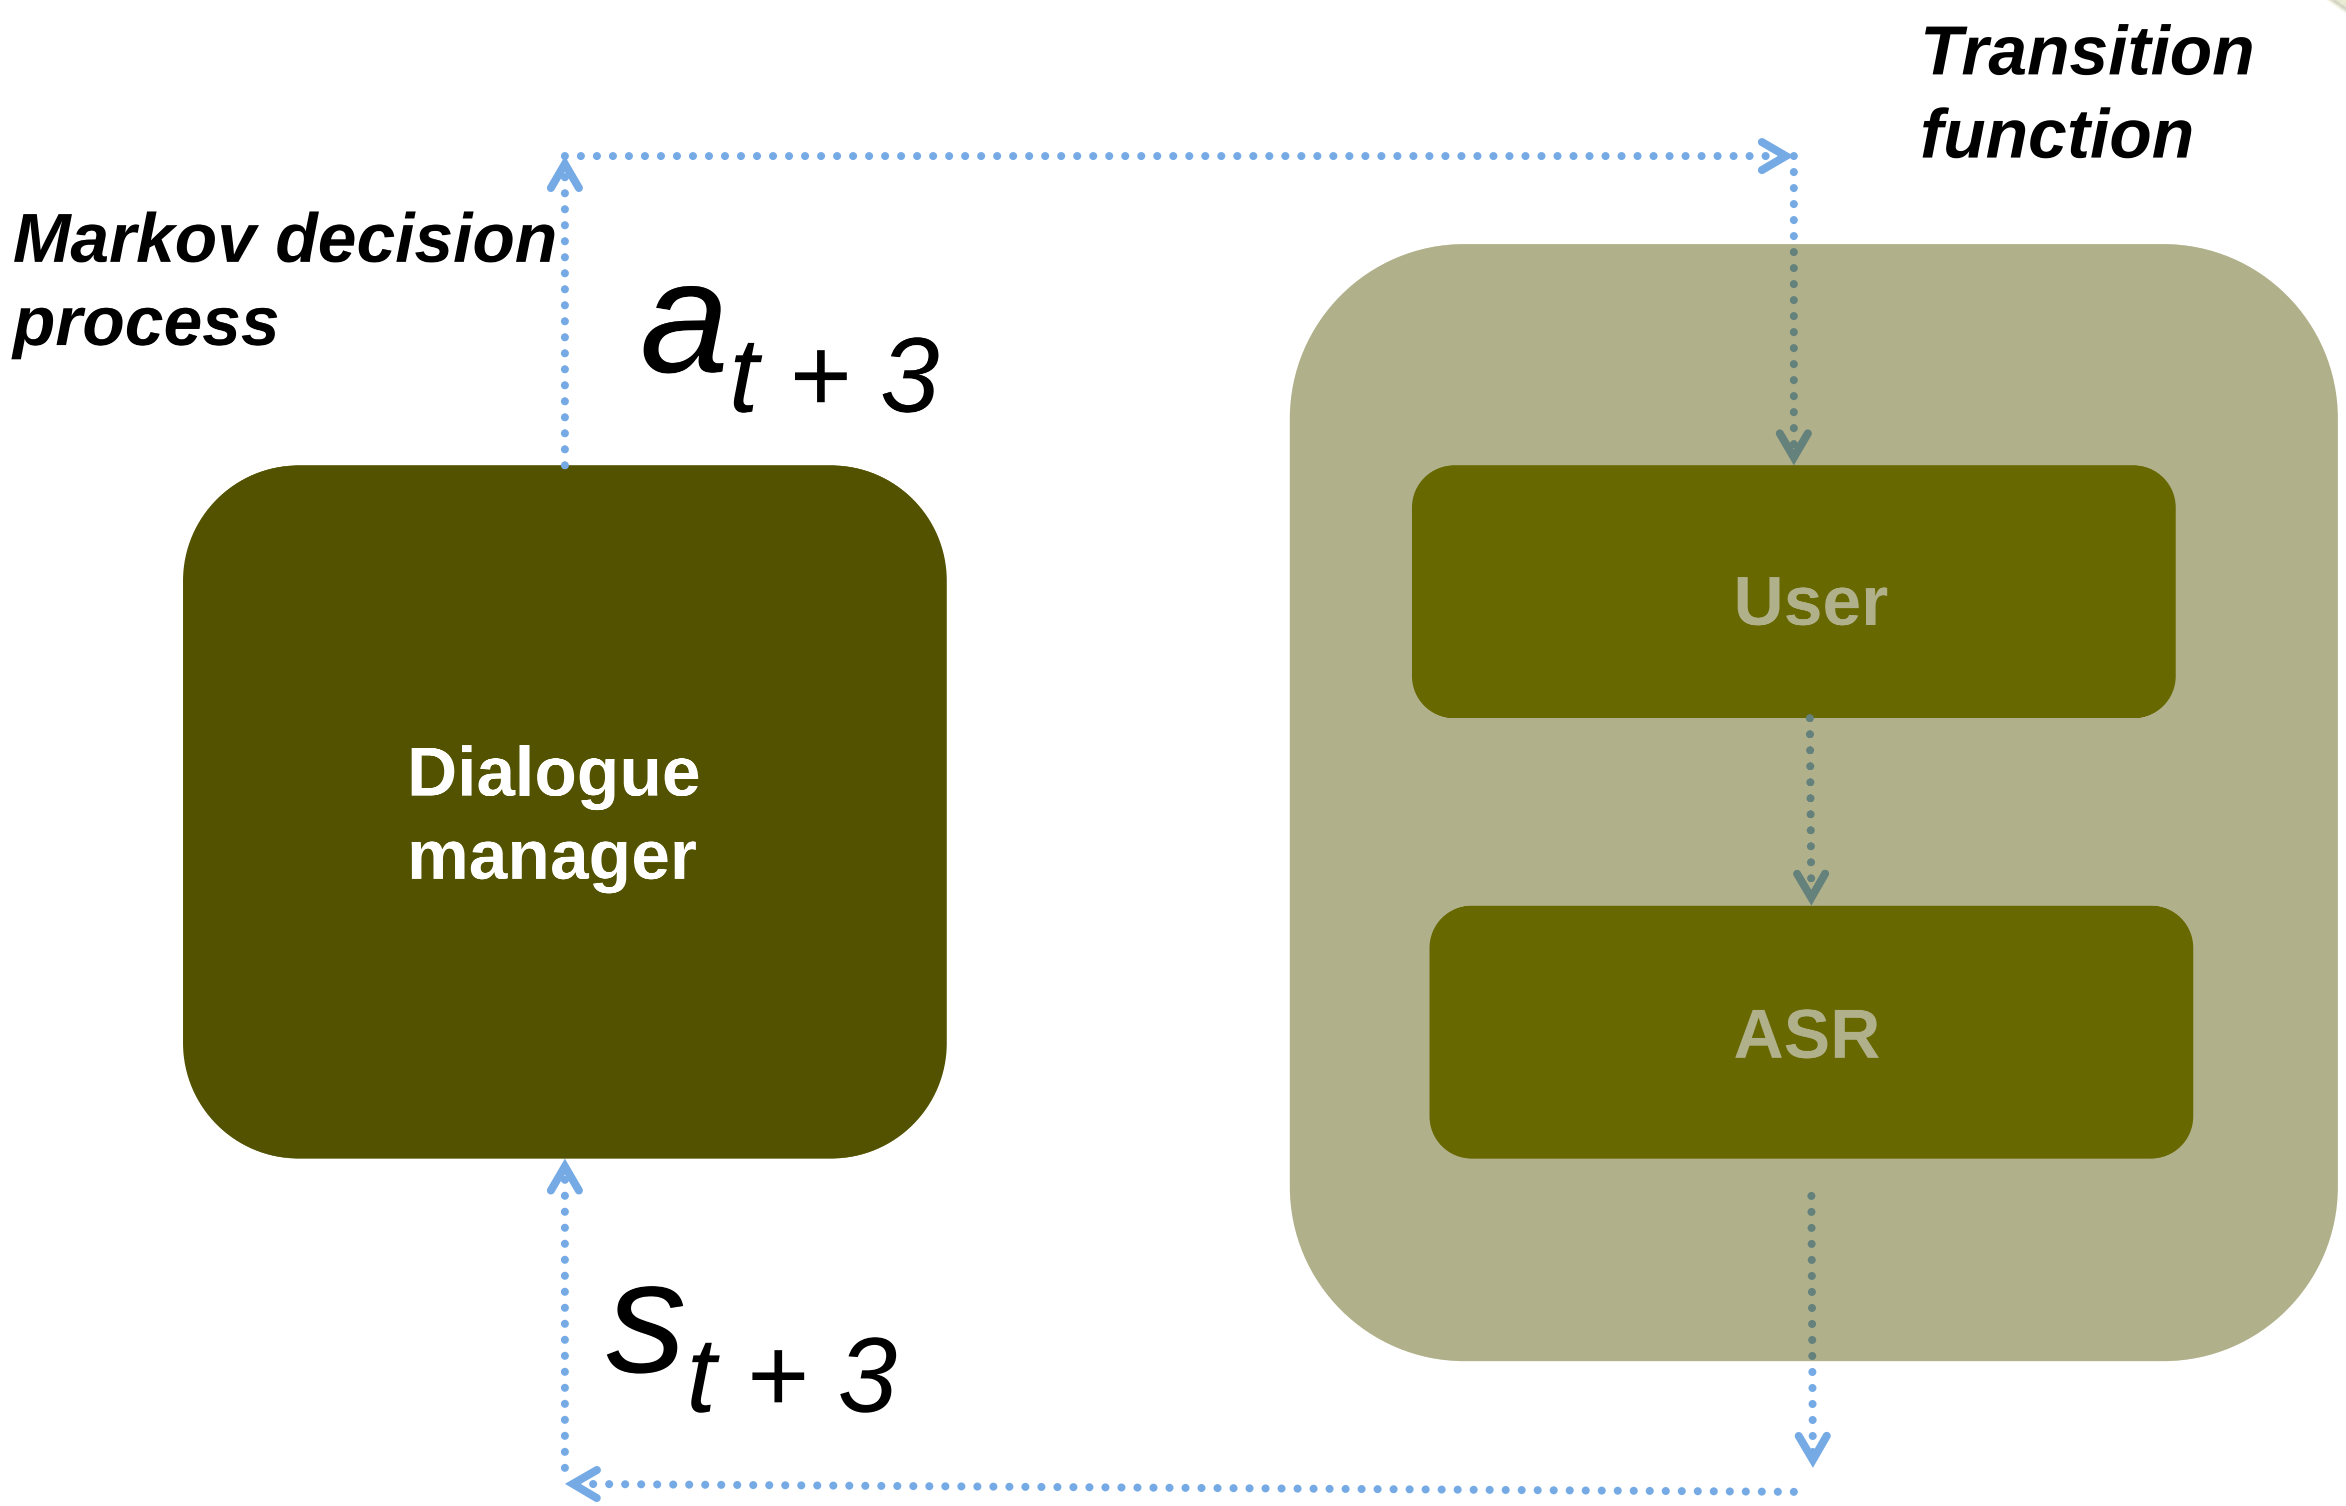

Intermezzo 2 - Learning

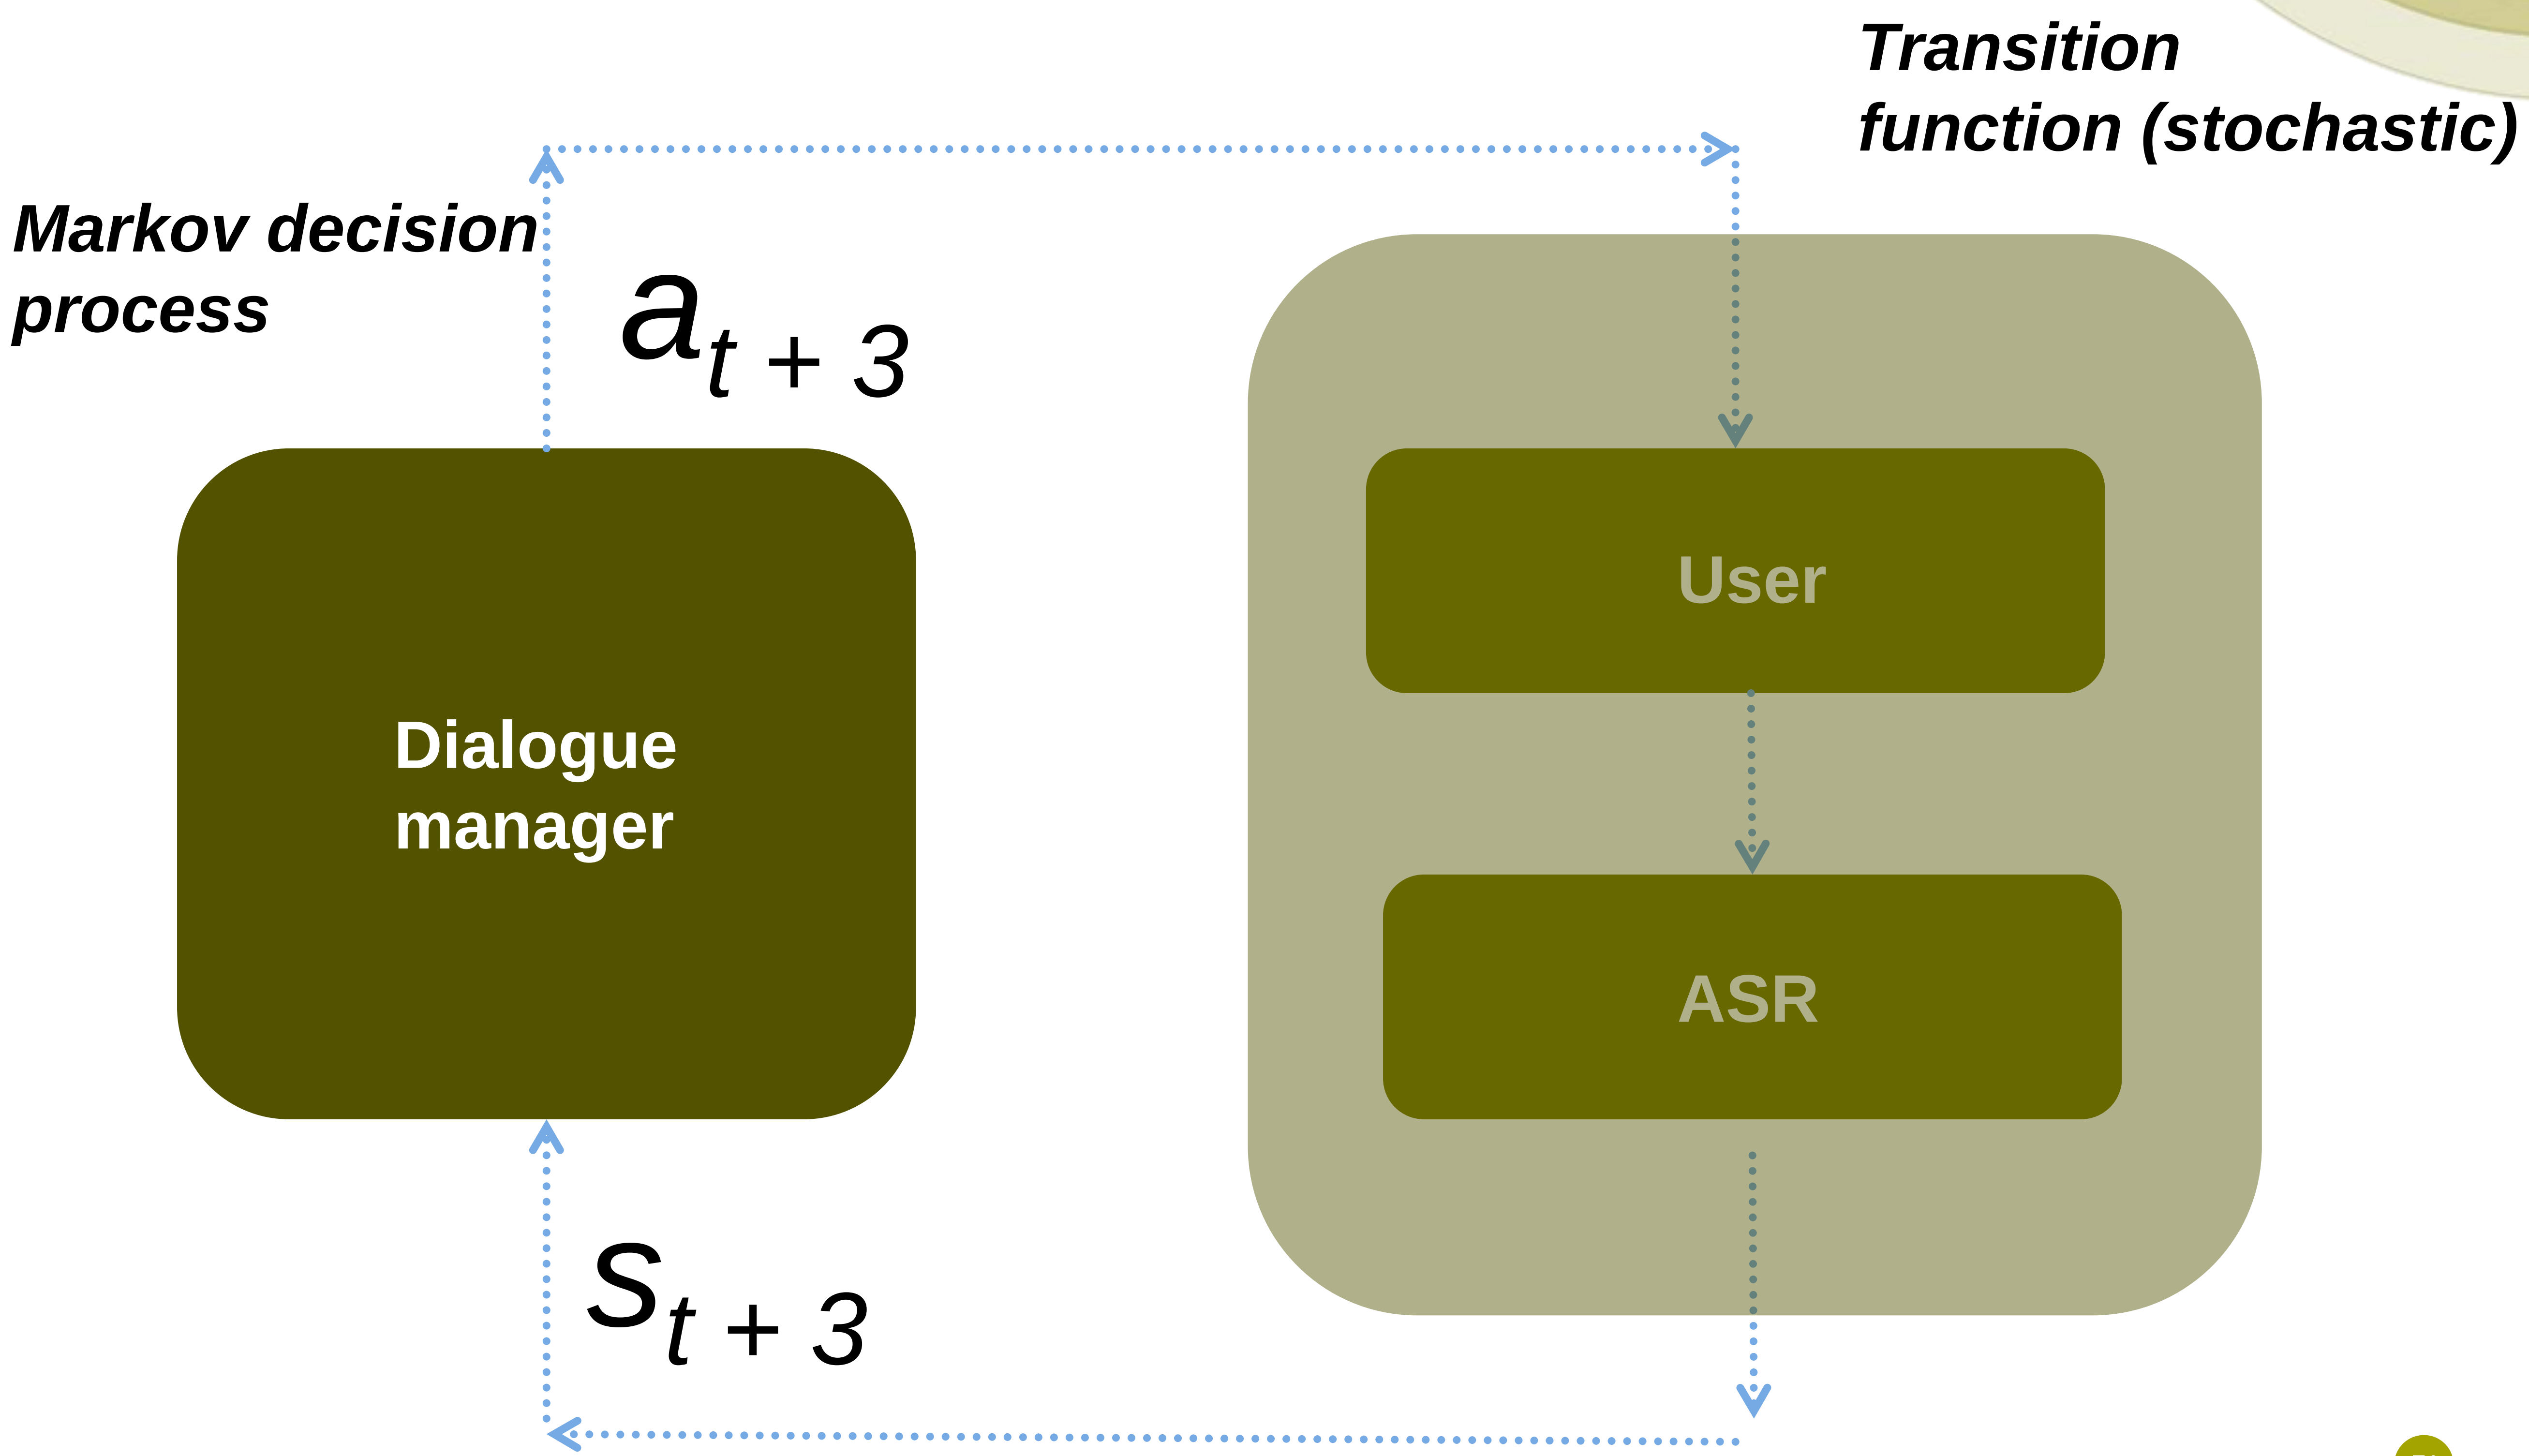

Intermezzo 2 - Learning

Reinforcement learning

- **Actions**: ask(x), confirm(x), askconf(x,y), askagain(y), accept(x)
- **Policy**:

	a1	a2	a3	...	an
s1	0.1	0.8	...	...	...
s2					
s3					
...					
sn					

- **Expected reward** of a policy in the final state.
- **Training** on real data or simulated transition function/user (select concept and apply noise)

Intermezzo 2 - Learning

Reinforcement learning – Young (2000: 1396)

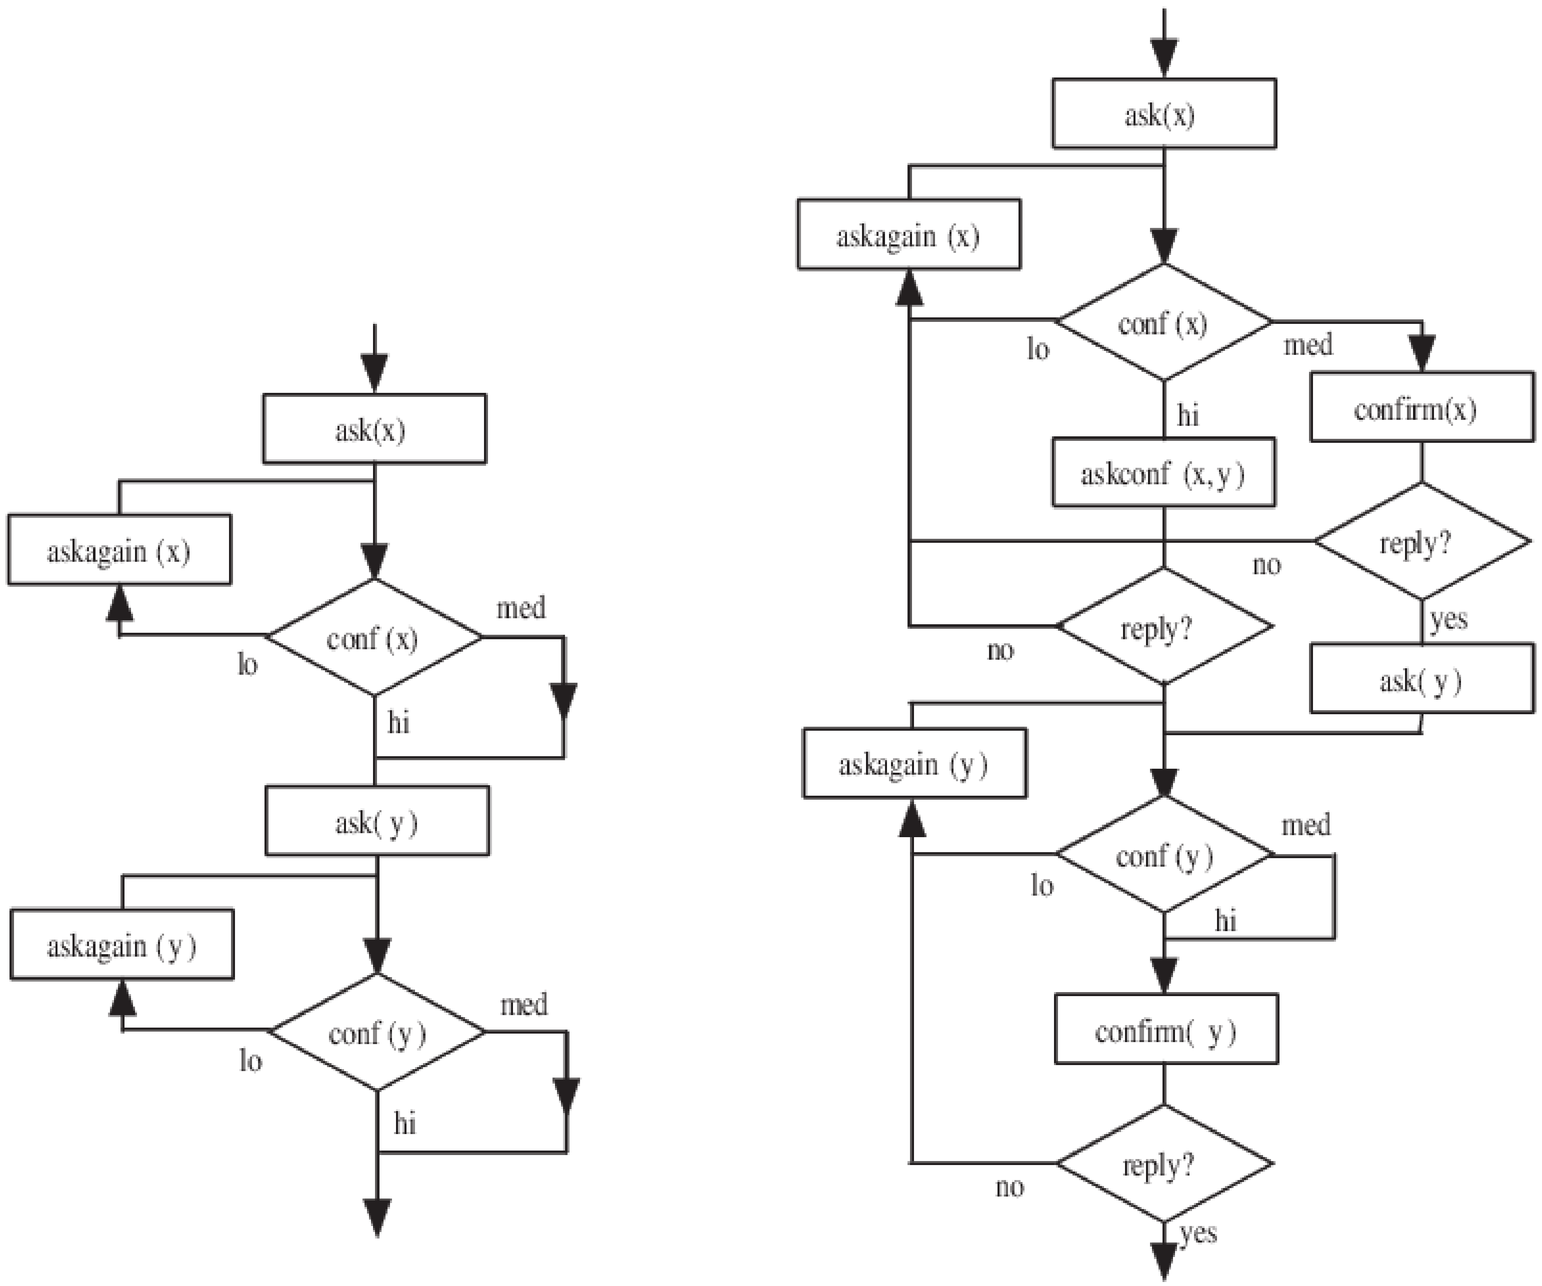

(a) Near-perfect recognition (b) Imperfect recognition

Young (Cambridge)
Georgila (ICT - USC)
Lemon (Heriot-Watt)
...

Intermezzo 3 – More evaluation

Dialogue State Tracking challenge, DSTC 2 (Smith, 2014)

Compare dialogue system representation of state with actual information.

System

- Area Chelsea
- Food Chinese
- Name -
- Price £10 - £40

Actual

- Area -
- Food French
- Name Exquis
- Price £10 - £40

Correct = 0
Extraneous attributes = 1
Missing attributes = 1
False attributes = 1

Selected-task agenda

- Personal assistants: Microsoft Cortana, Google Now and Apple Siri
- These systems can cope with a range of tasks:
  - launching an application
  - sending messages
  - accessing restaurant recommendations
  - adding reminders to the calendar
  - setting the alarm
  - searching on the web

Selected-task agenda

Siri (July, 2015)

1	User	*Activates Siri by pressing a specific button on the phone or saying ‘Hey Siri’*
2	Siri	What can I help you with?
3	User	I would like to make a call.
4	Siri	With whom would you like to speak?
5	User	Joe Bloggs
6	Siri	Just to confirm – you’d like to call Joe Bloggs?
		[Call]
7	User	*Selects call*
8	Siri	Calling Joe Bloggs

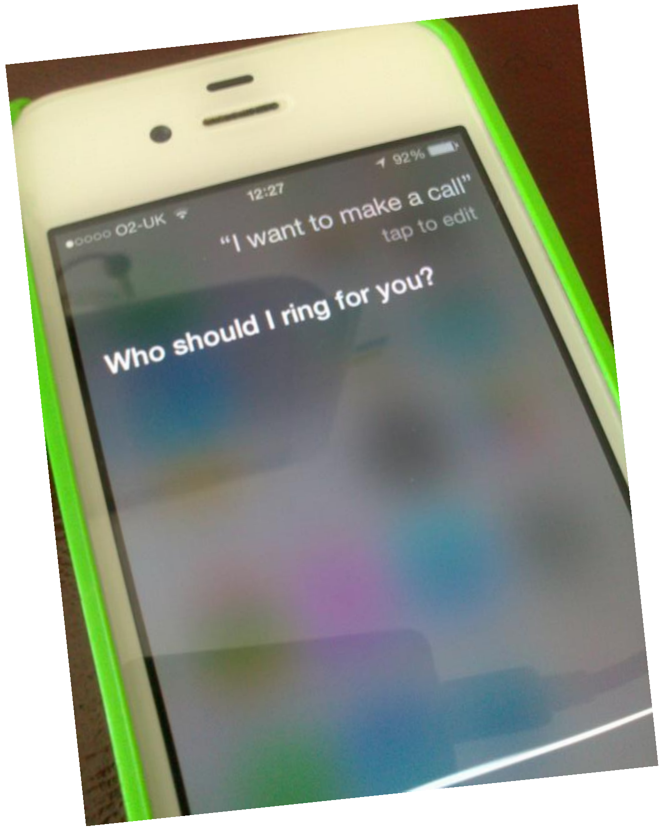

Selected-task agenda

The store

Agenda

Common

Selected-task agenda

The store

Agenda

Common

User I would like to make a call.

Selected-task agenda

The store

Agenda

1. Ask for callee

Common

Attribute	Value	Confidence
Callee		

Selected-task agenda

The store

Agenda

1. Ask for callee

Common

Attribute	Value	Confidence
Callee		

Siri With whom would you like to speak?

Select-task agenda

Mobile phone dialogue

- Access to various sources of contextual information
- Multimodality

Where is the sea?

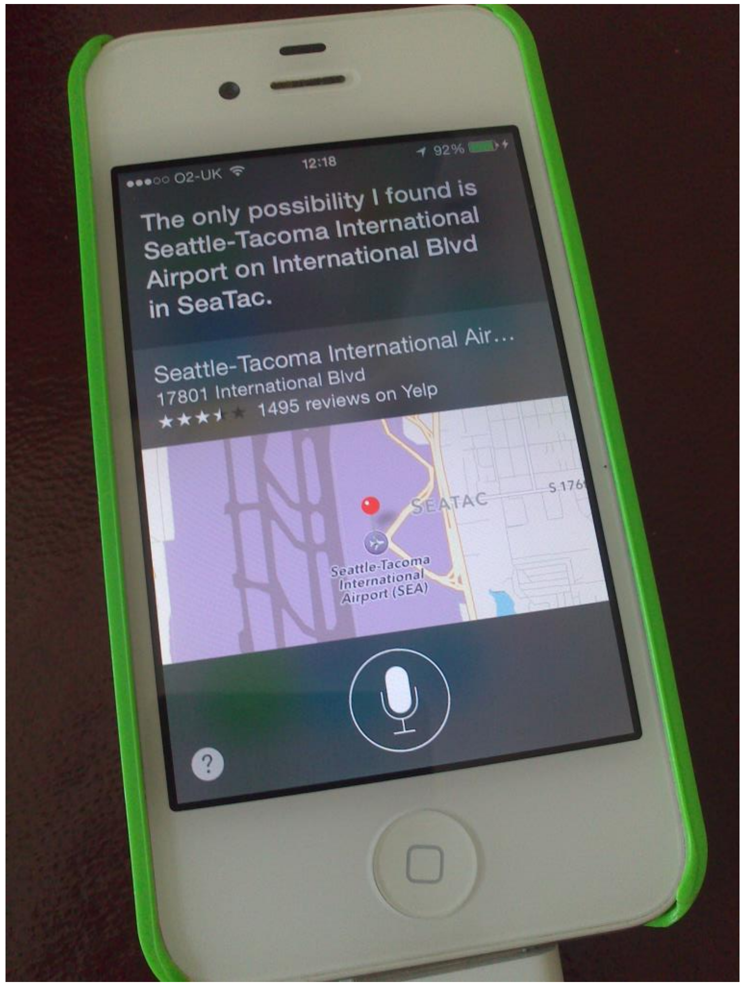

Joint-task agenda

- The task is negotiated as part of the conversation, rather than selected by one of the participants.
- Interlocutors collaborate on achieving this task.

Joint-task agenda

Power (1974, 1979: 109)

We are not yet able to construct formally precise theories of advanced cognitive processes such as language understanding

(…)

The present model will have served its purpose if it

(a) highlights some problems in the organisation of dialogue which the reader may not have explicitly noticed,
(b) explores a clear set of ideas for solving these problems (the most important idea in this case being the ‘conversational procedure’), and
(c) exposes the limitations of these ideas and therefore helps someone to construct a better theory.

The program

1. Data structures representing a **world** (with its own objects, laws, etc.) that the agents inhabit;
2. **Robots** (John and Mary) each with a *mind* and certain *capabilities*;
3. A “**chairman**” who arranges time sharing between the agents;
4. Functions that print out what is happening.

The program

World

- **Objects** = John, Mary, a door, (a bolt)
- **State** = position of objects:
  - For robots: IN/OUT
  - For the door: OPEN/SHUT
  - For the bolt: UP/DOWN
- **Actions**
  - For robots: MOVE
  - For the door: PUSH
  - For the bolt: SLIDE

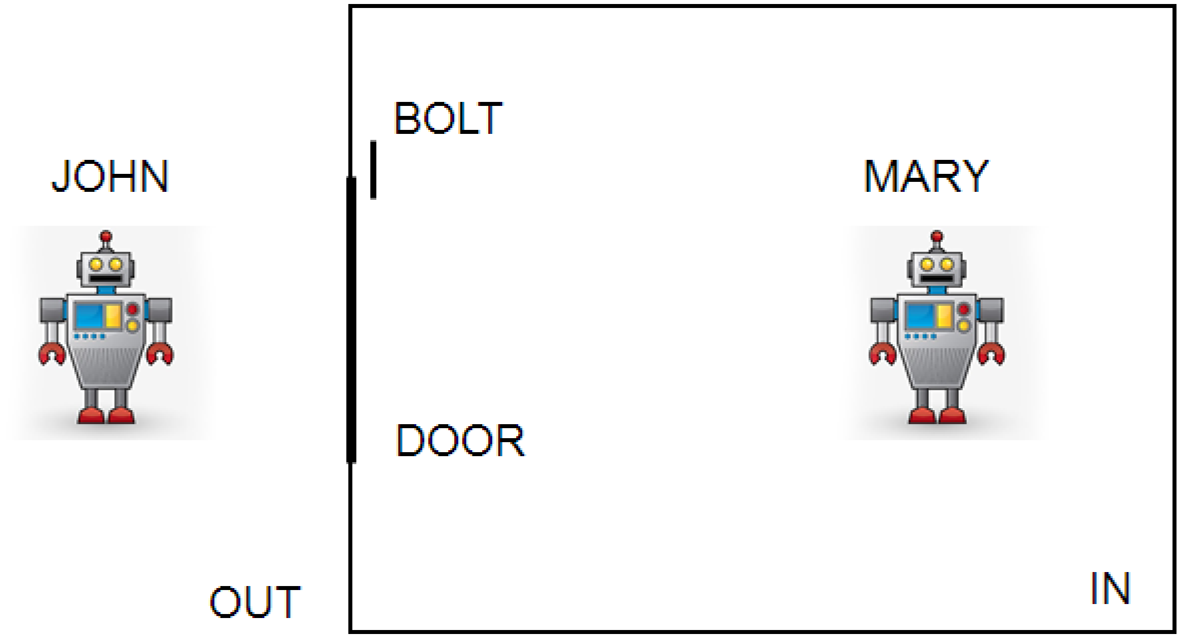

The program

World

“**Laws of nature**”:

If robot MOVE, then position robot changes provided that door OPEN.

If robot PUSH door, then position door changes provided that bolt UP.

If robot SLIDES bolt, then position bolt changes provided that robot is IN.

Program

Robots

Planning tree (populated with goals)
E.g.: John's goal: JOHN IN

Procedure execution stack

Beliefs (about regularities in the world)

John:

If a robot MOVES, nothing happens
If a robot PUSHES the door, the door changes position
If a robot SLIDES the bolt, nothing happens

Mary:

If a robot MOVES, it changes position provided the door is OPEN.
If a robot PUSHES the door, the door changes position.
If a robot SLIDES the bolt, nothing happens.

Program

Robots

Each robot has **capabilities** for:

Perception:

John: SEE
Mary: -

Action:

John: MOVE, SLIDE
Mary: MOVE, SLIDE, PUSH

Planning and Conversational Procedures

Conversation example

John	I want to suggest a goal.
Mary	Go ahead.
John	Will you help me get in?
Mary	By all means.
John	Shall we make a plan.
Mary	May I ask you something
John	Go ahead.
Mary	Are you in?
John	No.
Mary	Shall we make a plan?
John	Okay.
Mary	Is the door open?
John	No
Mary	I suggest that we get the door open and then you move.
John	I want to explain something.
Mary	Go ahead.
John	If one moves, nothing happens.
Mary	I disagree. If one moves when the door is open one changes position.
John	I see.
Mary	I suggest that we get the door open and then you move.
John	All right.
Mary	Shall we make a plan?
John	Okay.
Mary	I suggest that I push the door.

Conversation example

John **I want to suggest a goal.**
Mary **Go ahead.**
John **Will you help me get in?**
Mary **By all means.**
John Shall we make a plan.
Mary May I ask you something
John Go ahead.
Mary Are you in?
John No.
Mary Shall we make a plan?
John Okay.
Mary Is the door open?
John No
Mary I suggest that we get the door open and then you move.
John I want to explain something.
Mary Go ahead.
John If one moves, nothing happens.
Mary I disagree. If one moves when the door is open one changes position.
John I see.
Mary I suggest that we get the door open and then you move.
John All right.
Mary Shall we make a plan?
John Okay.
Mary I suggest that I push the door.

John in

current joint

Conversation example

John	**I want to suggest a goal.**
Mary	**Go ahead.**
John	**Will you help me get in?**
Mary	**By all means.**
John	Shall we make a plan.
Mary	May I ask you something
John	Go ahead.
Mary	Are you in?
John	No.
Mary	Shall we make a plan?
John	Okay.
Mary	Is the door open?
John	No
Mary	**I suggest that we get the door open and then you move.**
John	I want to explain something.
Mary	Go ahead.
John	If one moves, nothing happens.
Mary	I disagree. If one moves when the door is open one changes position.
John	I see.
Mary	I suggest that we get the door open and then you move.
John	All right.
Mary	Shall we make a plan?
John	Okay.
Mary	I suggest that I push the door.

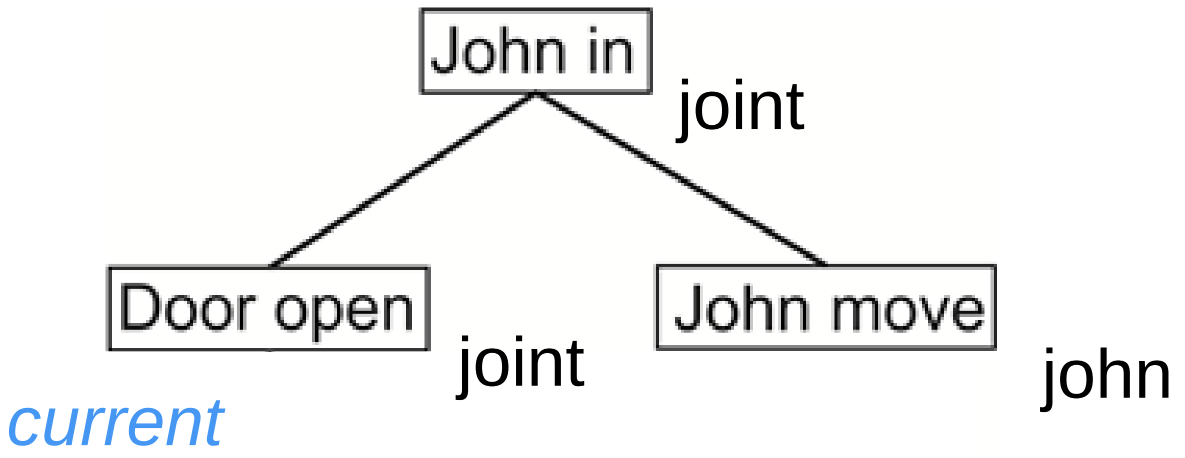

Conversation example

John	**I want to suggest a goal.**
Mary	**Go ahead.**
John	**Will you help me get in?**
Mary	**By all means.**
John	Shall we make a plan.
Mary	May I ask you something
John	Go ahead.
Mary	Are you in?
John	No.
Mary	Shall we make a plan?
John	Okay.
Mary	Is the door open?
John	No
Mary	**I suggest that we get the door open and then you move.**
John	I want to explain something.
Mary	Go ahead.
John	If one moves, nothing happens.
Mary	I disagree. If one moves when the door is open one changes position.
John	I see.
Mary	I suggest that we get the door open and then you move.
John	All right.
Mary	**Shall we make a plan?**
John	**Okay.**
Mary	**I suggest that I push the door.**

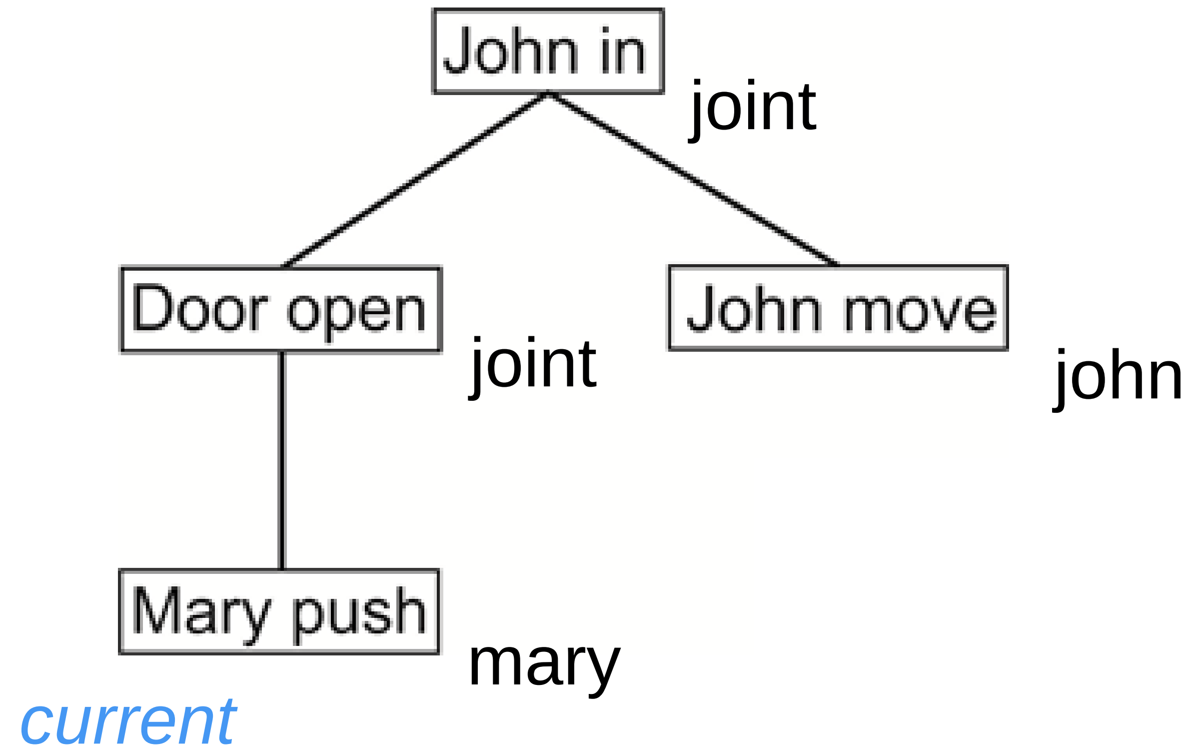

Procedures

CHOOSEGOAL

1. Select a main goal and call it M
2. Try to achieve main goal M on your own (individual ACHIEVEGOAL).
3. If result of ACHIEVEGOAL is succeed, go to 5.
4. Secure cooperation through AGREEGOAL. If answer A "yes", label M as a joint responsibility and call joint ACHIEVEGOAL. If the answer A is "no", go to 5.
5. do nothing (control back to the chairman)

Procedures

ACHIEVEGOAL

1. Identify current goal G
2. If action G and can be performed directly then
    1. if G is assigned to partner, wait till they say they've performed it;
    2. if G is assigned to you, perform it. If the goal above G is joint, inform your partner that you've performed it and initiate a subconversation to assess the result.

   Remove G from the tree and return to 1.
3. If state G, test whether it holds. If so, remove it. If main goal, exit and report success; else return to 1.
4. Else, FINDPLAN P to achieve G. If there is no plan, remove G and its sister(s) and record failure, and go to 1. If G is main, return *fail*.
5. Attach P to G and return to 0.

Procedures

FINDPLAN

1. Identify the type of object in goal G (e.g, for JOHN IN, ROBOT)
2. Find an applicable belief **B** (e.g. <u>If</u> action MOVE, <u>then</u> change ROBOT IN/OUT <u>provided that</u> DOOR OPEN) Report failure if no belief could be found.
3. If the goal is joint, set ROBOT to the robot who can perform the plan to achieve the goal. If neither can, exit with *fail*.
4. Else if the goal is your responsibility, check that you can do the action and set ROBOT to you. If you can't do the action, exit with *fail*.
5. Replace ROBOT in **B** with the designated agent.
6. Test whether the "provided that" clause Z is satisfied.
7. If Z is already satisfied, P := X, else, P:= Z, X

Procedure

AGREEGOAL (S1,S2,G) ⇒ A

1. S1 composes sentence S asking for help with goal G
2. S2 decodes S, obtaining a value for G. If it is identical with their own main goal, they give A the value “yes”, if not, “no”. If A is “yes”, the variable A of CHOOSEGOAL is also set to “yes”. S2 utters A and exits from AGREEGOAL
3. S1 reads A and exits returning A to the procedure which called AGREEGOAL.

Power (1979)

Summary

- A program for simulating dialogues between agents;
- Based on traditional AI planning notions (goal, belief, plan, …);
- Tight integration of conversational and other planning procedures, e.g, ACHIEVEGOAL can call AGREEGOAL;
- A dialogue store with several sections:
  - Plan tree
  - Procedure stack
  - Beliefs

Power (1979)

Summary

By tracing back the procedure calls that gave rise to an utterance, we can identify the point of the utterance:

John and I are cooperating to achieve the goal John IN. This goal is not yet achieved and we're trying to agree a plan. I am to propose a plan and John is to evaluate it. I'm trying to find out whether I need to include opening the door in my plan.

John	I want to suggest a goal.
Mary	Go ahead.
John	Will you help me get in?
Mary	By all means.
John	Shall we make a plan.
Mary	May I ask you something
John	Go ahead.
Mary	Are you in?
John	No.
Mary	Shall we make a plan?
John	Okay.
Mary	Is the door open?

Power (1979)

Limitation

Problematic utterances:

When I said I was out I was joking/lying/...

Actually, I changed my mind. I don't want to get in.

"The problem with the actual program is that although it uses plans, candidate plans, and statements about whether goals have been achieved, these are not explicitly marked and [related]." (Power, 1979: 150)

John	I want to suggest a goal.
Mary	Go ahead.
John	Will you help me get in?
Mary	By all means.
John	Shall we make a plan.
Mary	May I ask you something
John	Go ahead.
Mary	Are you in?
John	No.
Mary	Shall we make a plan?
John	Okay.
Mary	Is the door open?

Non-cooperation

Plüss, Piwek, Power (2011), Plüss (2014)

Paxman	Are you proud of having got rid of one of the very few black women in Parliament?
Galloway	I'm not err... Jeremy, move on to your next question.
Paxman	You're not answering that one?
Galloway	No, because I don't believe that people get elected because of the colour of their skin. I believe people get elected because of their record and because of their policies. So move on to your next question.
Paxman	Are you proud...
Galloway	Because I've got a lot of people who want to speak to me.
Paxman	You...
Galloway	If you ask that question again, I'm going, I warn you now.
Paxman	Don't try and threaten me Mr Galloway, please.

Non-cooperation

Plüss, Piwek, Power (2011), Plüss (2014)

- If we go back to the notion of a dialogue game, we distinguished:
  - legitimate moves (as defined by the generation rules) – these determine the interlocutors' discourse obligations
  - a strategy for selecting an actual dialogue act, given the set of legitimate acts

- Non-cooperation:
  - Add special rules for non-cooperative dialogue, or
  - non-cooperative behaviour occurs when participants favour individual goals that are in conflict with their current discourse obligations.

Incremental dialogue processing

(Howes, Purver, Healey, Mills & Gregoromichelaki, 2011)

- Compound contributions: single syntactic or semantic unit which is divided over 2 or more contributions. Contributions are bounded by: change in speaker, significant pause, or end of a sentence. (20% in corpus study using portion of the BNC)

- Cross-speaker:

Daughter: Oh here dad, a good way to get those corners out
Dad: is to stick yer finger inside.
Daughter: well, that's one way.

Incremental dialogue processing

(DeVault, Sagae, Traum, 2009)

- We can provide you with power generators.
- Semantics:

$$\begin{bmatrix} mood : declarative \\ sem : \begin{bmatrix} type : event \\ agent : captain - kirk \\ event : deliver \\ theme : power - generator \\ modal : \begin{bmatrix} possibility : can \end{bmatrix} \\ speech - act : \begin{bmatrix} type : offer \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

- ASR build up representations for 1 word input, 2 word in put, etc. ASR trained on partial inputs.
- Machine learning to predict when further information doesn't improve interpretation

Incremental dialogue processing

(DeVault, Sagae, Traum, 2009)

- “elder do you agree to move the”
- Complete utterance by retrieving closes match from the corpus and presenting remaining words.
- “elder do you agree to move the clinic downtown?”
- Open question: When to generate such completions.

Summary

- **Dialogue Games**
- **Dialogue Systems**
  - *Reactive*
  - *Agenda-driven*
    - Fixed-task
    - Selected-task
    - Evaluation and learning
    - Joint-task
- **Open challenges**
  - *Non-cooperation*
  - *Incrementality*

Questions