UiO : University of Oslo

Dialogue Management with Probabilistic Rules

Pierre Lison
Language Technology Group
University of Oslo

November 10, 2014

Outline

- The dialogue management task
- Probabilistic rules
  - General idea
  - Parameter estimation
- Evaluation
- Conclusion

Dialogue architecture

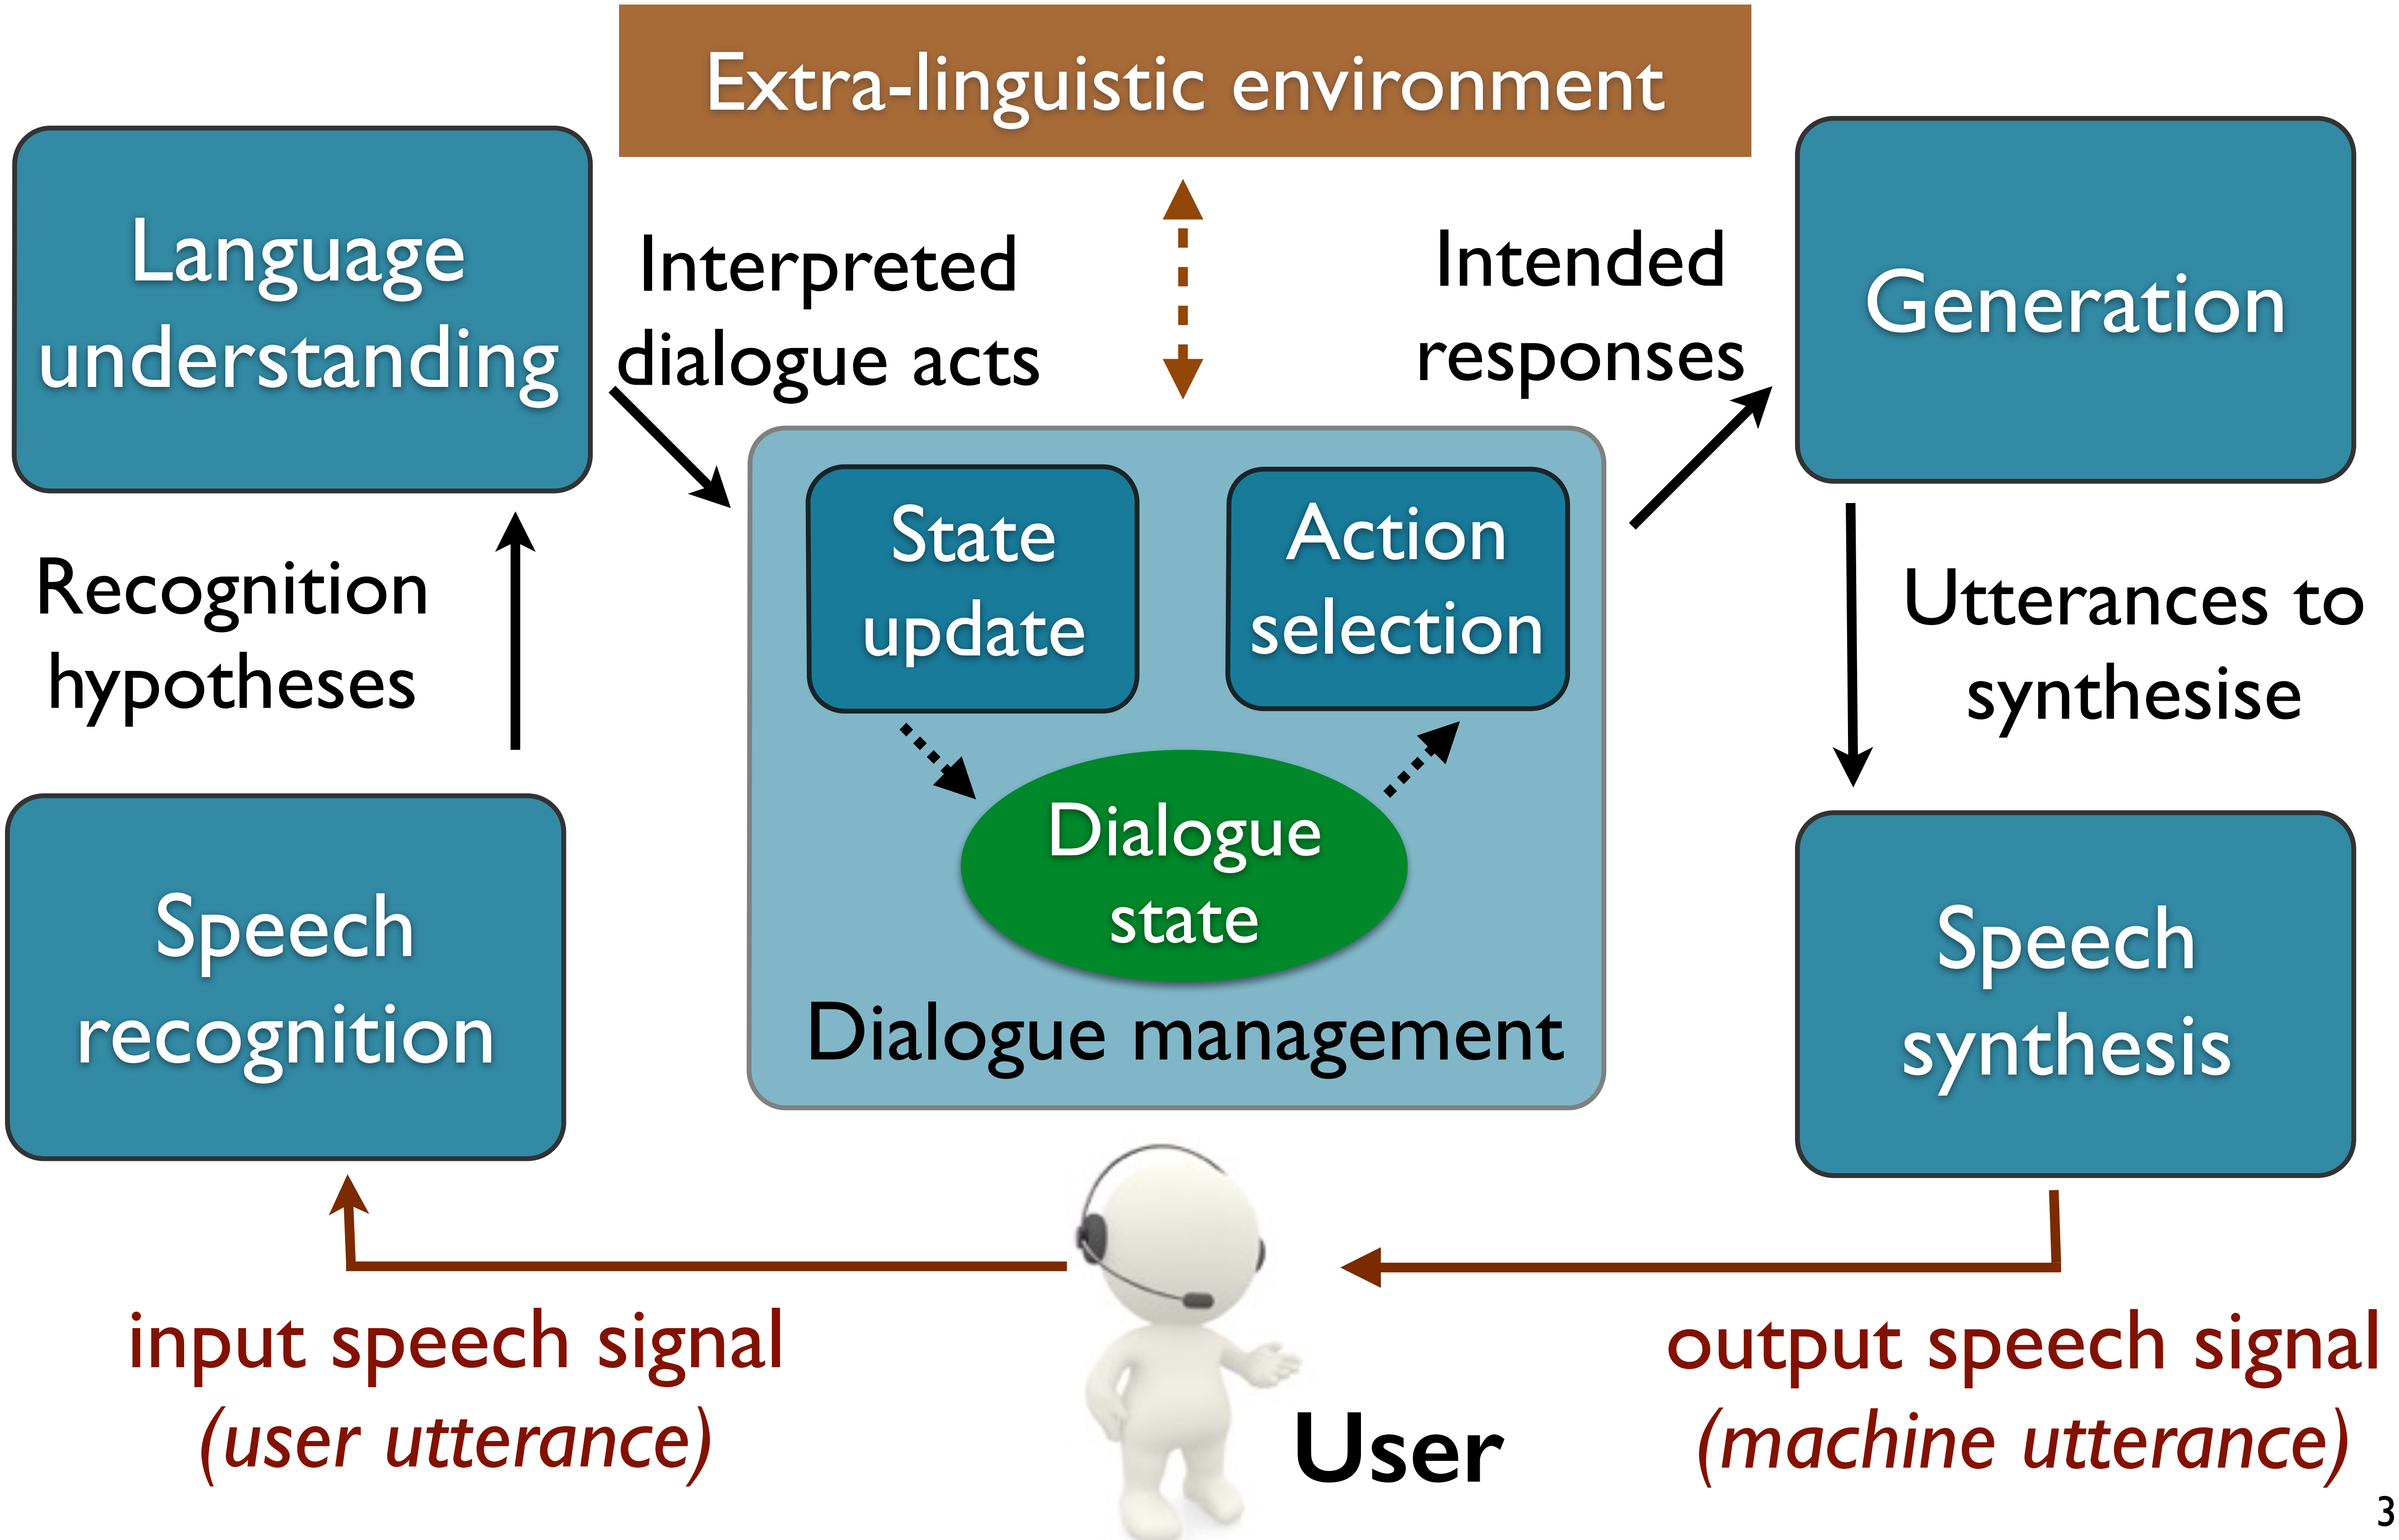

Two core challenges

Spoken dialogue is ...

Complex

- Context is essential to make sense of most dialogues
- Linguistic *and* extra-linguistic factors

Uncertain

- Pervasiveness of noise, errors and ambiguity
- Numerous sources of variability

Existing techniques

	Logical approaches	**Statistical approaches**
+	Fine-grained control of conversation	Robust, data-driven models of dialogue
-	Limited account for uncertainties	Need large quantities of training data

A new, hybrid modelling framework based on *probabilistic rules*

The approach

- Dialogue state encoded as a *Bayesian network*
  - Each state variable captures some relevant aspect of the interaction (dialogue history, user intentions, external environment, etc.)
  - The state is regularly updated upon new system actions and observations
  - ...And used to derive high-utility actions to execute

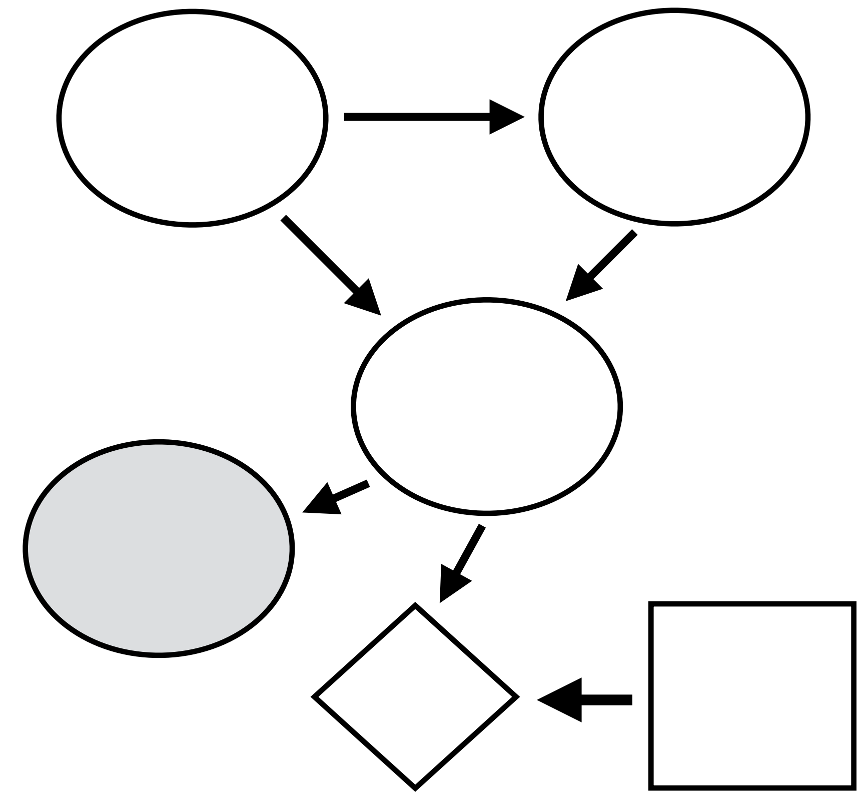

Similar to existing (PO)MDP approaches!

The approach

- **<u>But</u>**: instead of expressing the domain models using traditional formats, we adopt a high-level representation based on **probabilistic rules**

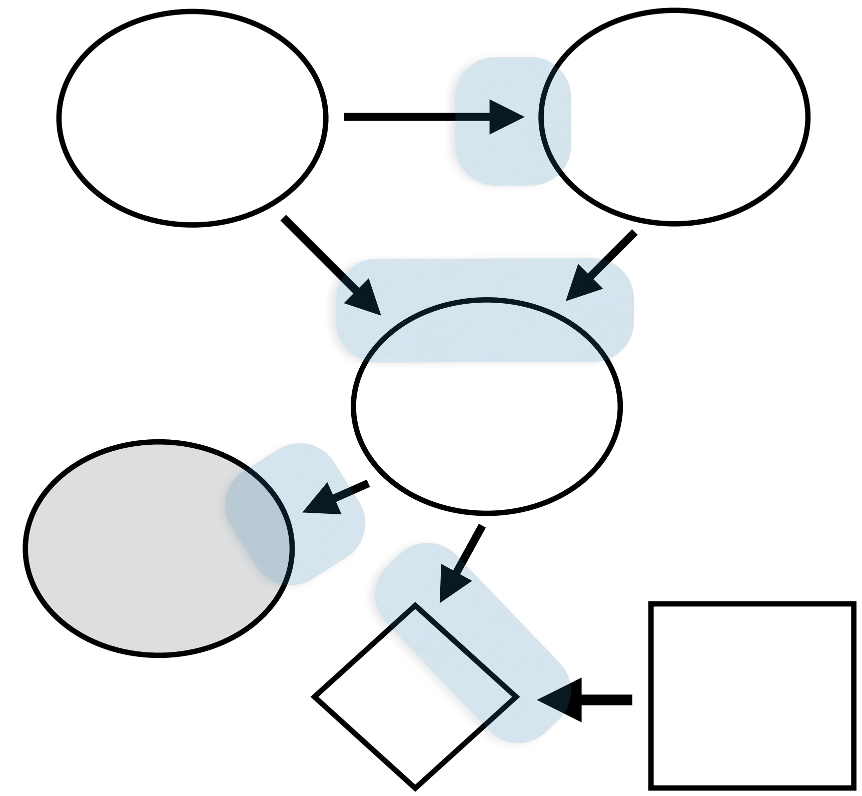

- Two main advantages:
  - Reduce the numbers of *unknown parameters* → easier to learn from limited amounts of data
  - Easier to integrate *expert knowledge* (in human-readable form)

Types of rules

	Probability rules	**Utility rules**
What they encode:	*Conditional probability distributions* between state variables	*Utility functions* for system actions given state variables
General skeleton:	**if** (condition$_1$) **then** $P(\text{effect}_1) = \theta_1$, $P(\text{effect}_2) = \theta_2$, ... **else if** (condition$_2$) **then** $P(\text{effect}_3) = \theta_3$, ...	**if** (condition$_1$) **then** $U(\text{action}_1) = \theta_1$, $U(\text{action}_2) = \theta_2$, ... **else if** (condition$_2$) **then** $U(\text{action}_3) = \theta_3$,... ...

Example of probability rule

$\forall$ x,

if (*last-user-input* = *x* ∧ *system-action* = AskRepeat) **then**

P(*next-user-input* = *x*) = 0.9

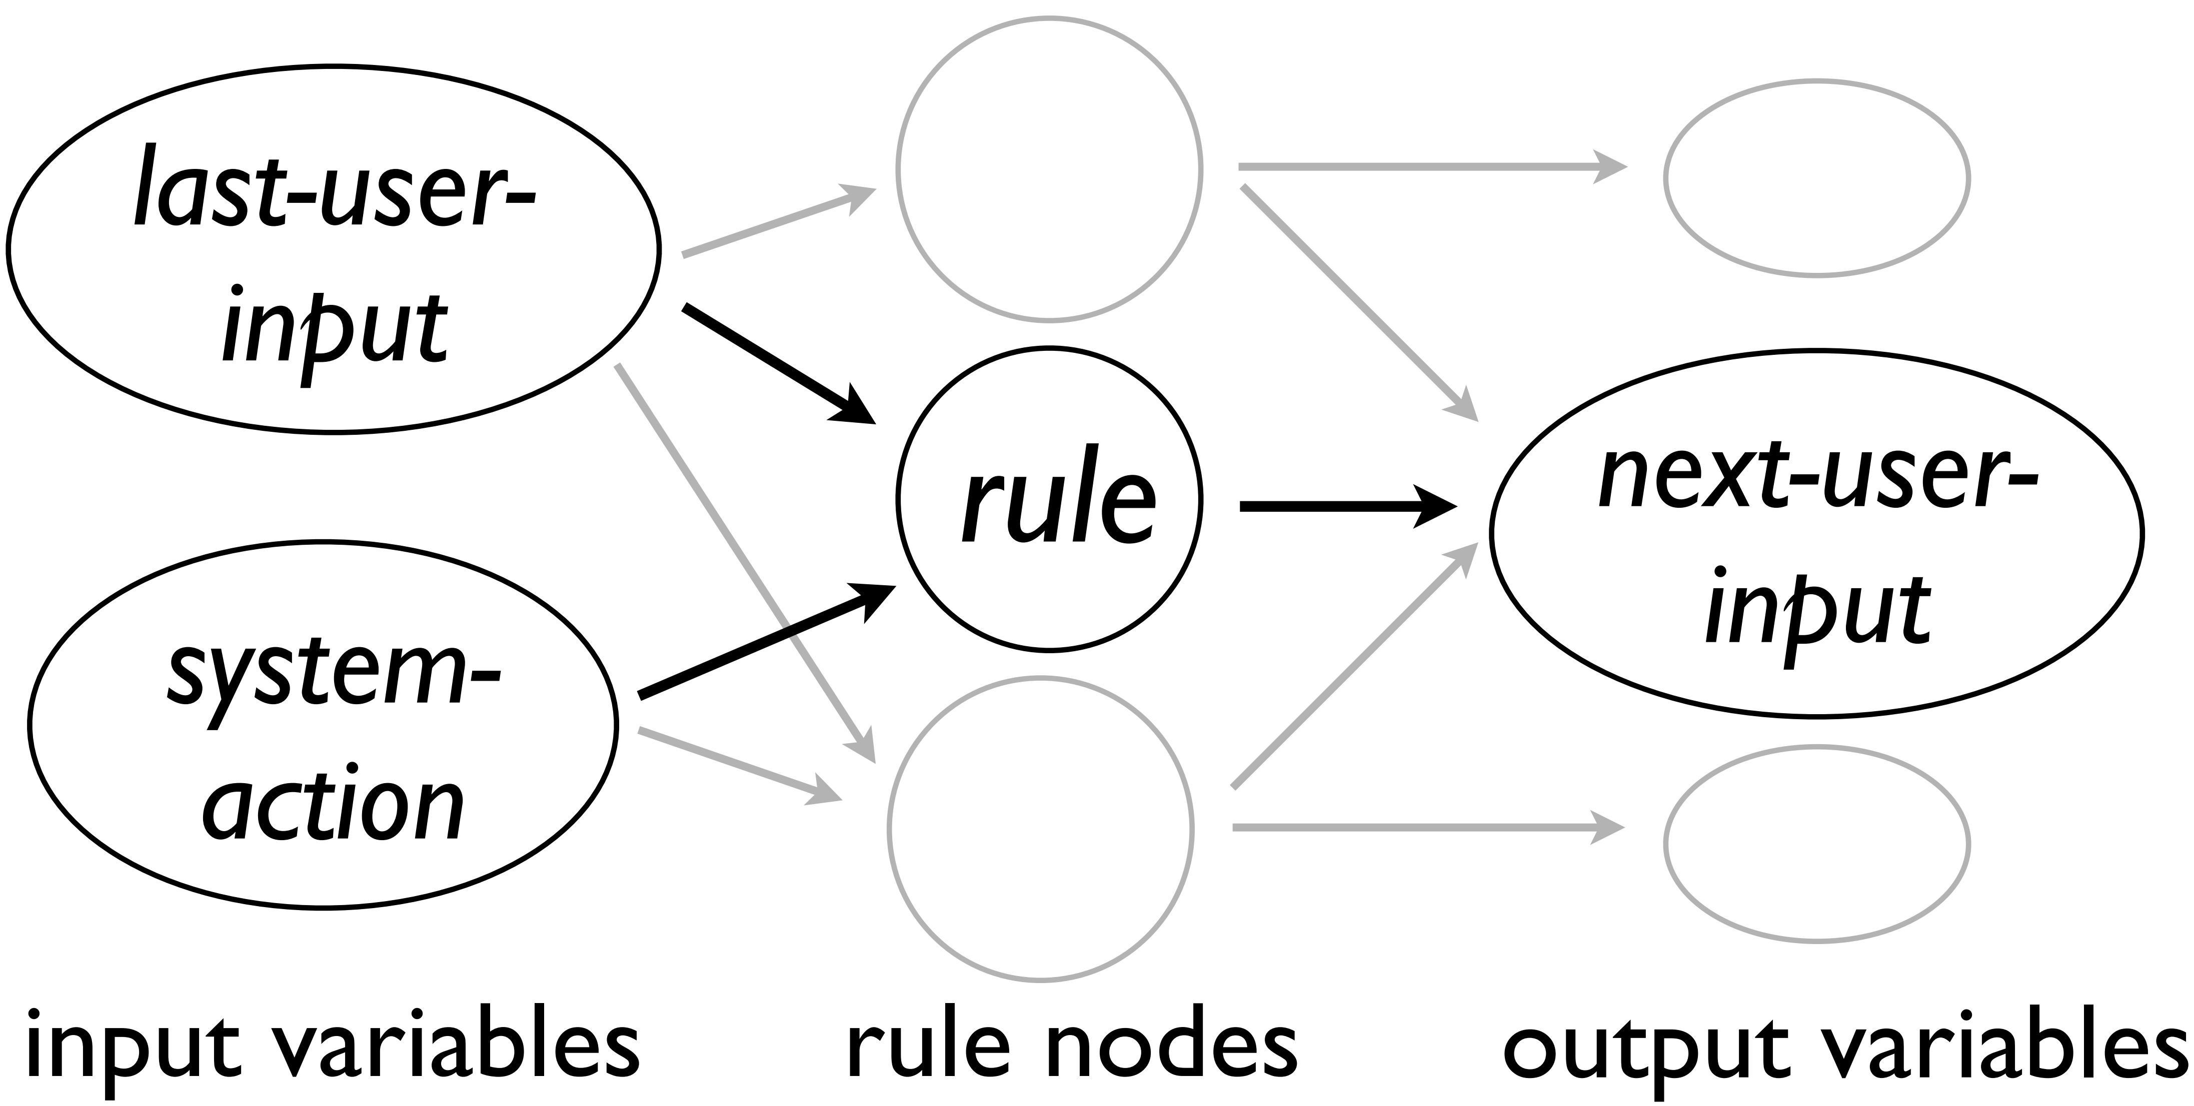

Example of utility rule

∀ x,

if (*last-user-input*=Request(x) ∧ x ∈ *perceived-objects*) **then**

U(*system-action*=PickUp(x)) = +5

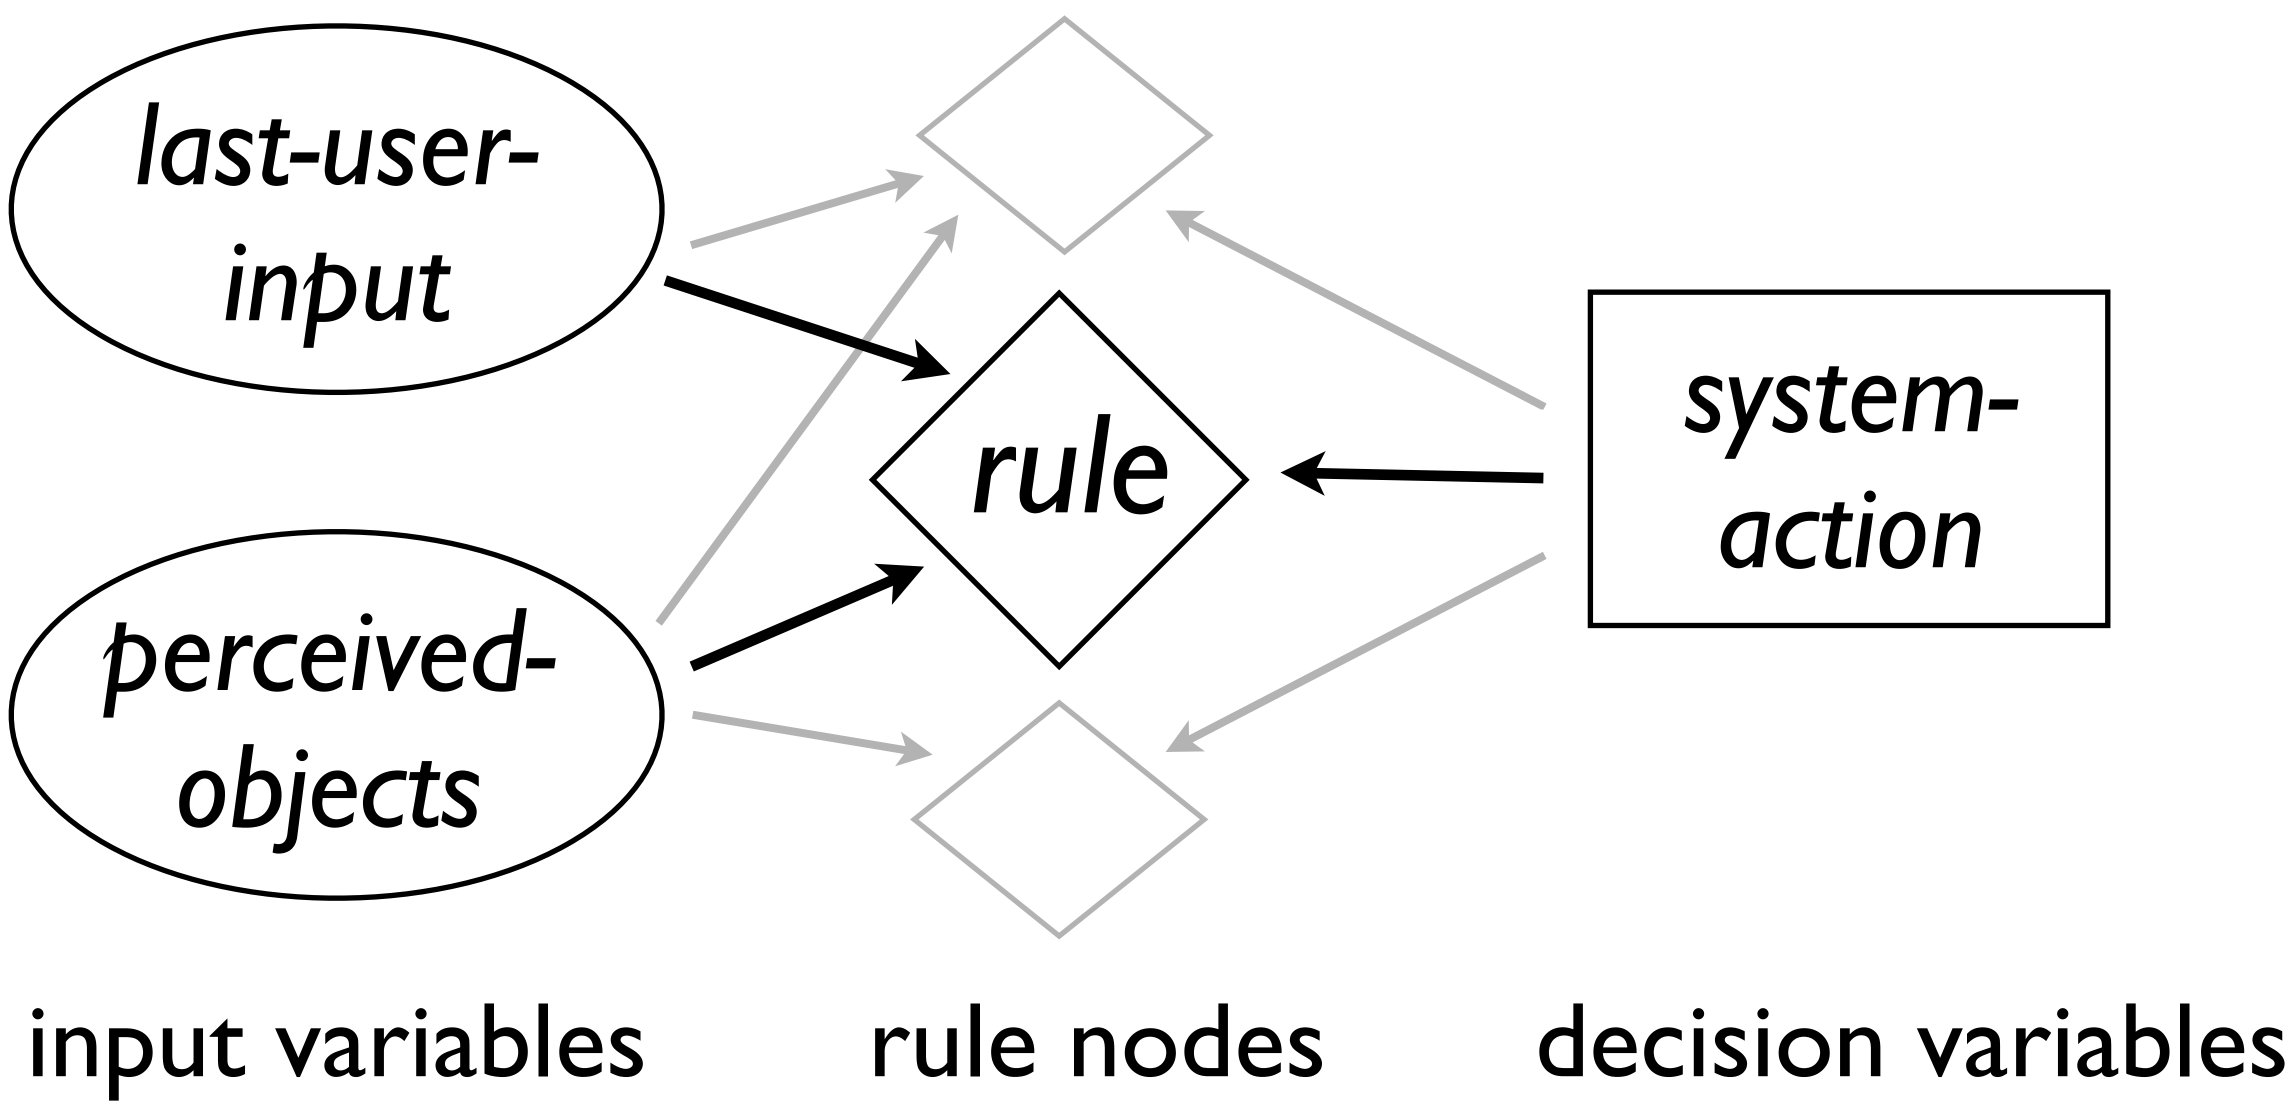

Processing workflow

- Dialogue state expressed as a Bayesian network
- External modules add new observations
- Probability rules used to update the dialogue state
- Utility rules used to select the system actions
- Implementation in the OpenDial toolkit [http://opendial.googlecode.com]

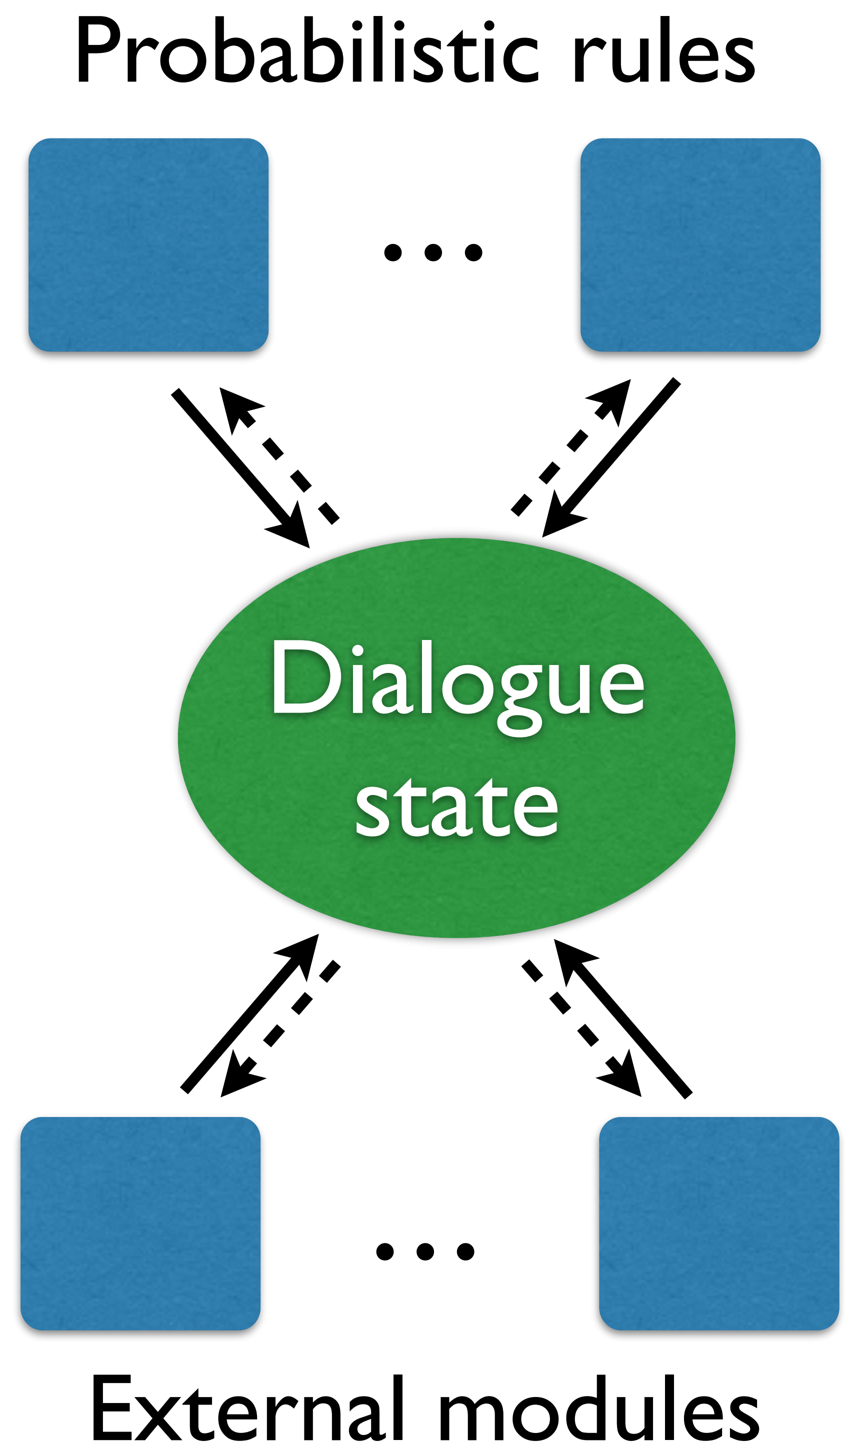

Parameter estimation

- Probabilistic rules may include parameters (unknown probabilities or utilities)
- **Bayesian learning** approach:
  - Start with initial prior over possible parameter values
  - Refine the distribution given the observed data $\mathcal{D}$

$$\underbrace{P(\boldsymbol{\theta} \,|\, \mathcal{D})}_{\text{Posterior distribution}} = \underbrace{\eta}_{\text{Normalisation factor}} \; \underbrace{P(\mathcal{D}\,;\boldsymbol{\theta})}_{\text{Likelihood of the data}} \; \underbrace{P(\boldsymbol{\theta})}_{\text{Prior distribution}}$$

Parameter estimation

$\forall$ x,

if (*last-user-input* = x $\wedge$ *system-action* = AskRepeat) **then**

P(*next-user-input* = x) = θ

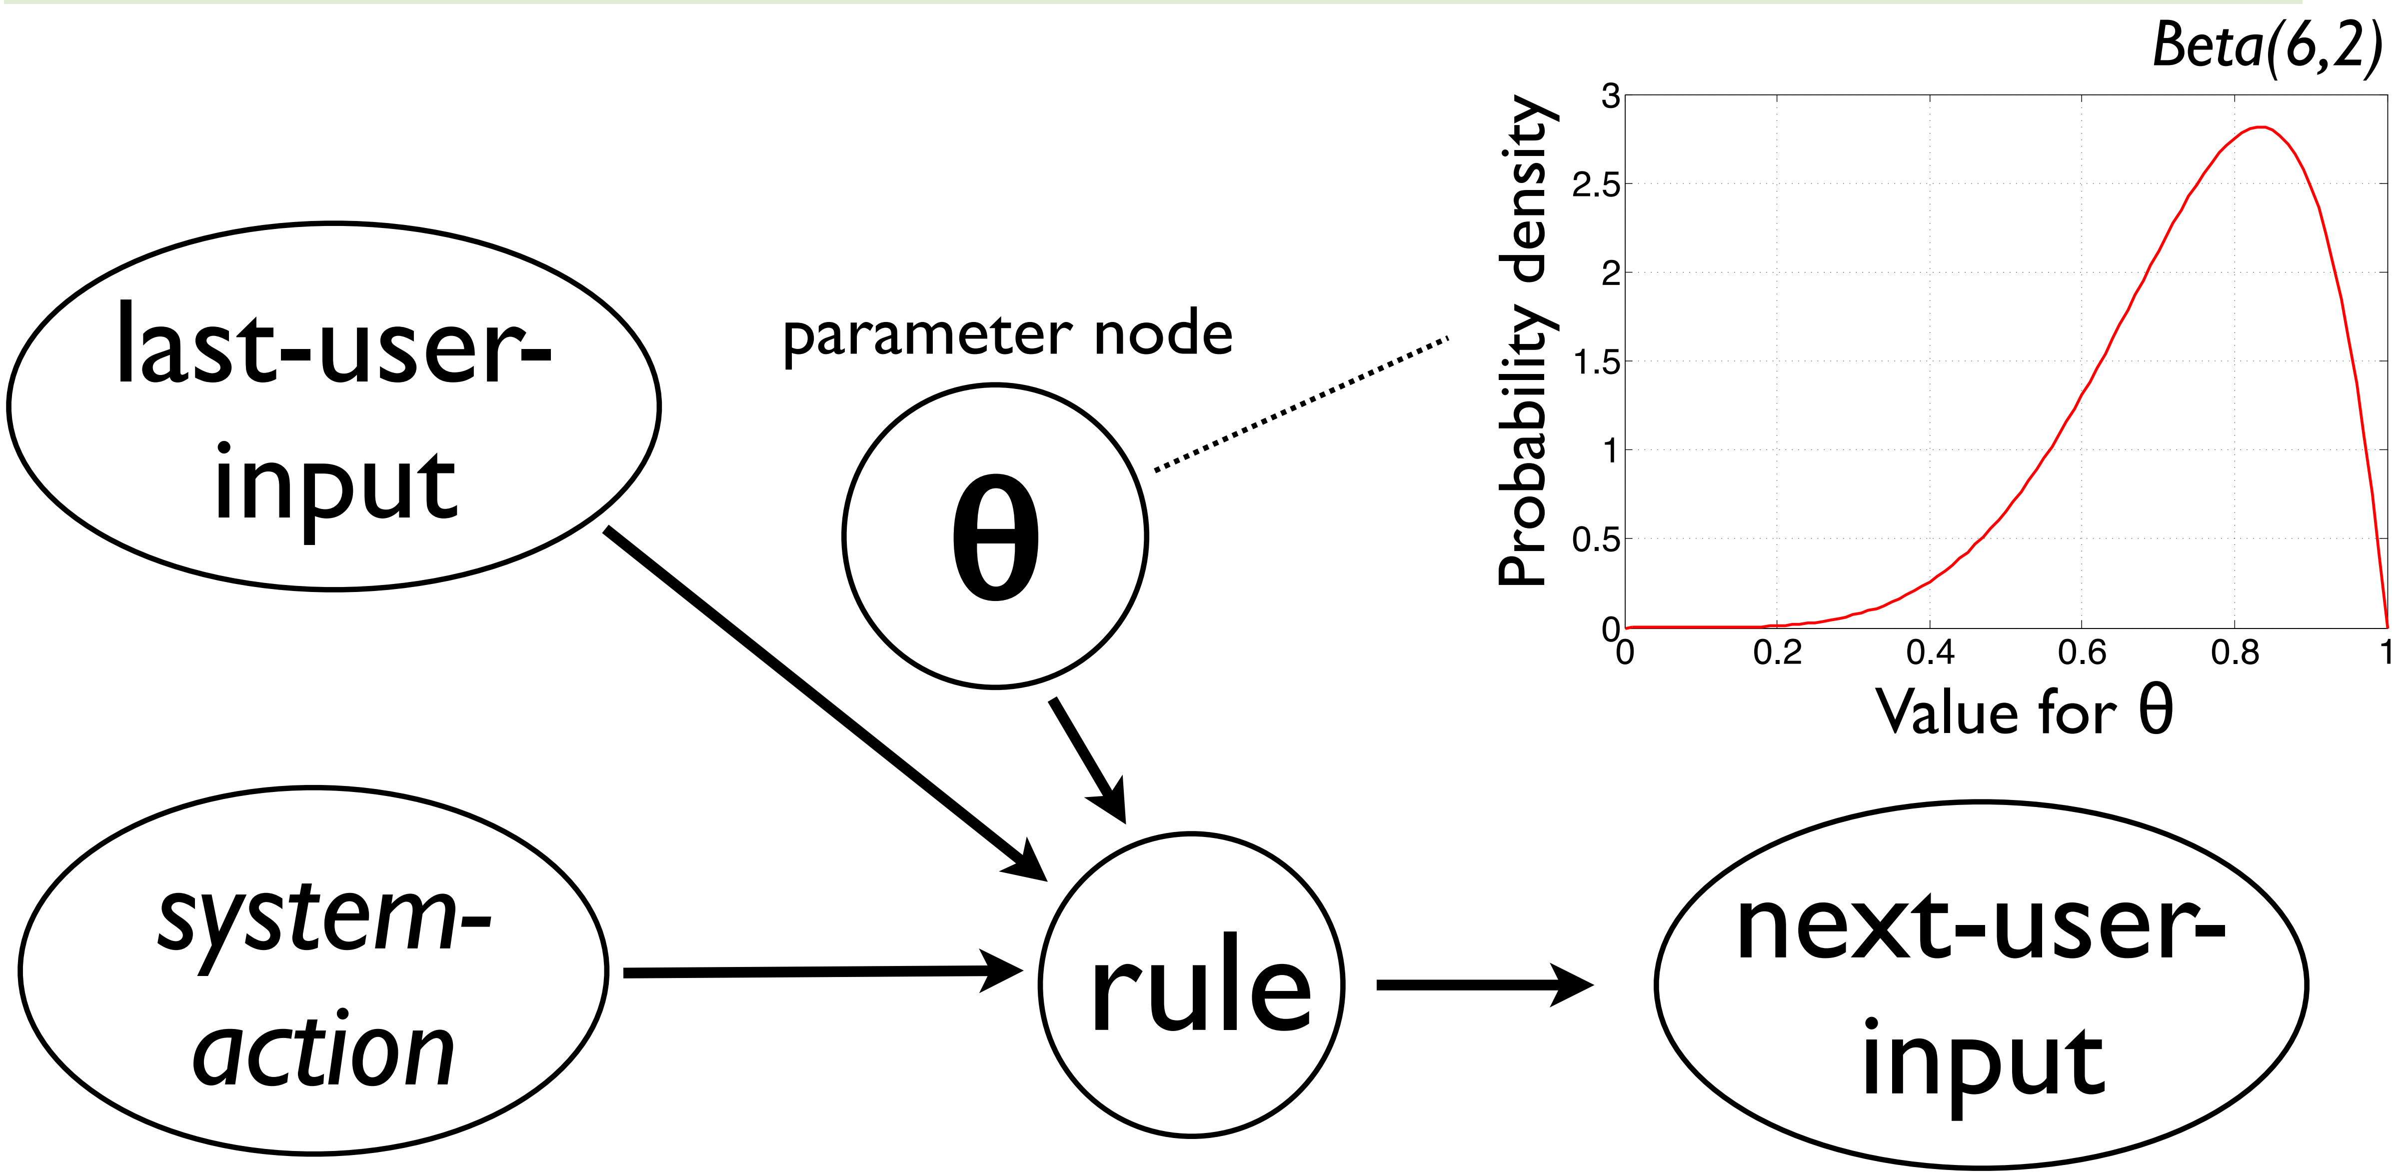

Learning paradigms

- Different types of training data:
  - *Supervised learning*: Wizard-of-Oz interactions

Goal: find the parameter values that best "imitate" the Wizard's conversational behaviour

  - *Reinforcement learning*: real or simulated interactions

Goal: find the parameter values that provide the best fit for the collected observations

[P. Lison. Probabilistic Dialogue Models with Prior Domain Knowledge (SIGDIAL 2012)]
[P. Lison. Model-based Bayesian Reinforcement Learning for Dialogue Management (Interspeech 2013)]

User evaluation

- **Task**: instruct the robot to move across the table, pick one cylinder and release it on the landmark

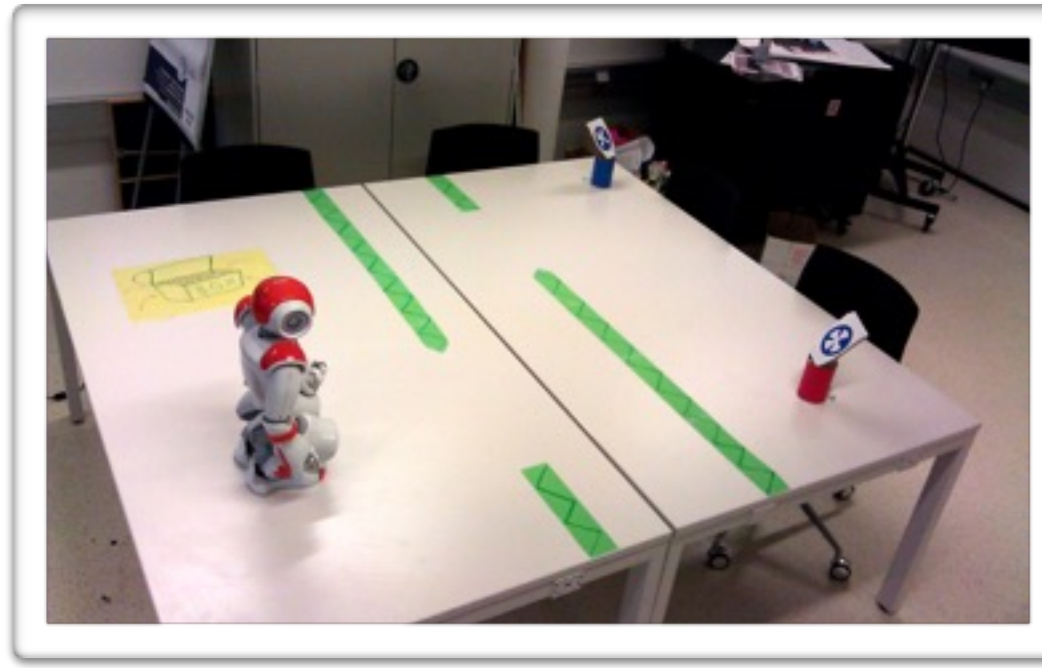

- Comparison of three modelling approaches:

1. A handcrafted finite-state automaton
2. A factored statistical model
3. A model structured with probabilistic rules

Experimental procedure

- **Step 1:** collect Wizard-of-Oz interaction data
- **Step 2:** Estimate the internal parameters for the 3 models with the collected data
- **Step 3:** Conduct user trials for the 3 approaches
- **Step 4:** Compare them on dialogue quality metrics

Dialogue domain:

- 26 user actions
- 41 system actions
- State size: 35×10^6 (10 variables)

Parameter estimation:

- 10 recorded WoZ interactions
- 3 parameters in handcrafted automaton (thresholds)
- 433 parameters in factored statistical model
- 28 parameters in model encoded with probabilistic rules

Learning curve

Training: 9 Wizard-of-Oz interactions (770 system turns)
Testing: 1 Wizard-of-Oz interaction (71 system turns)

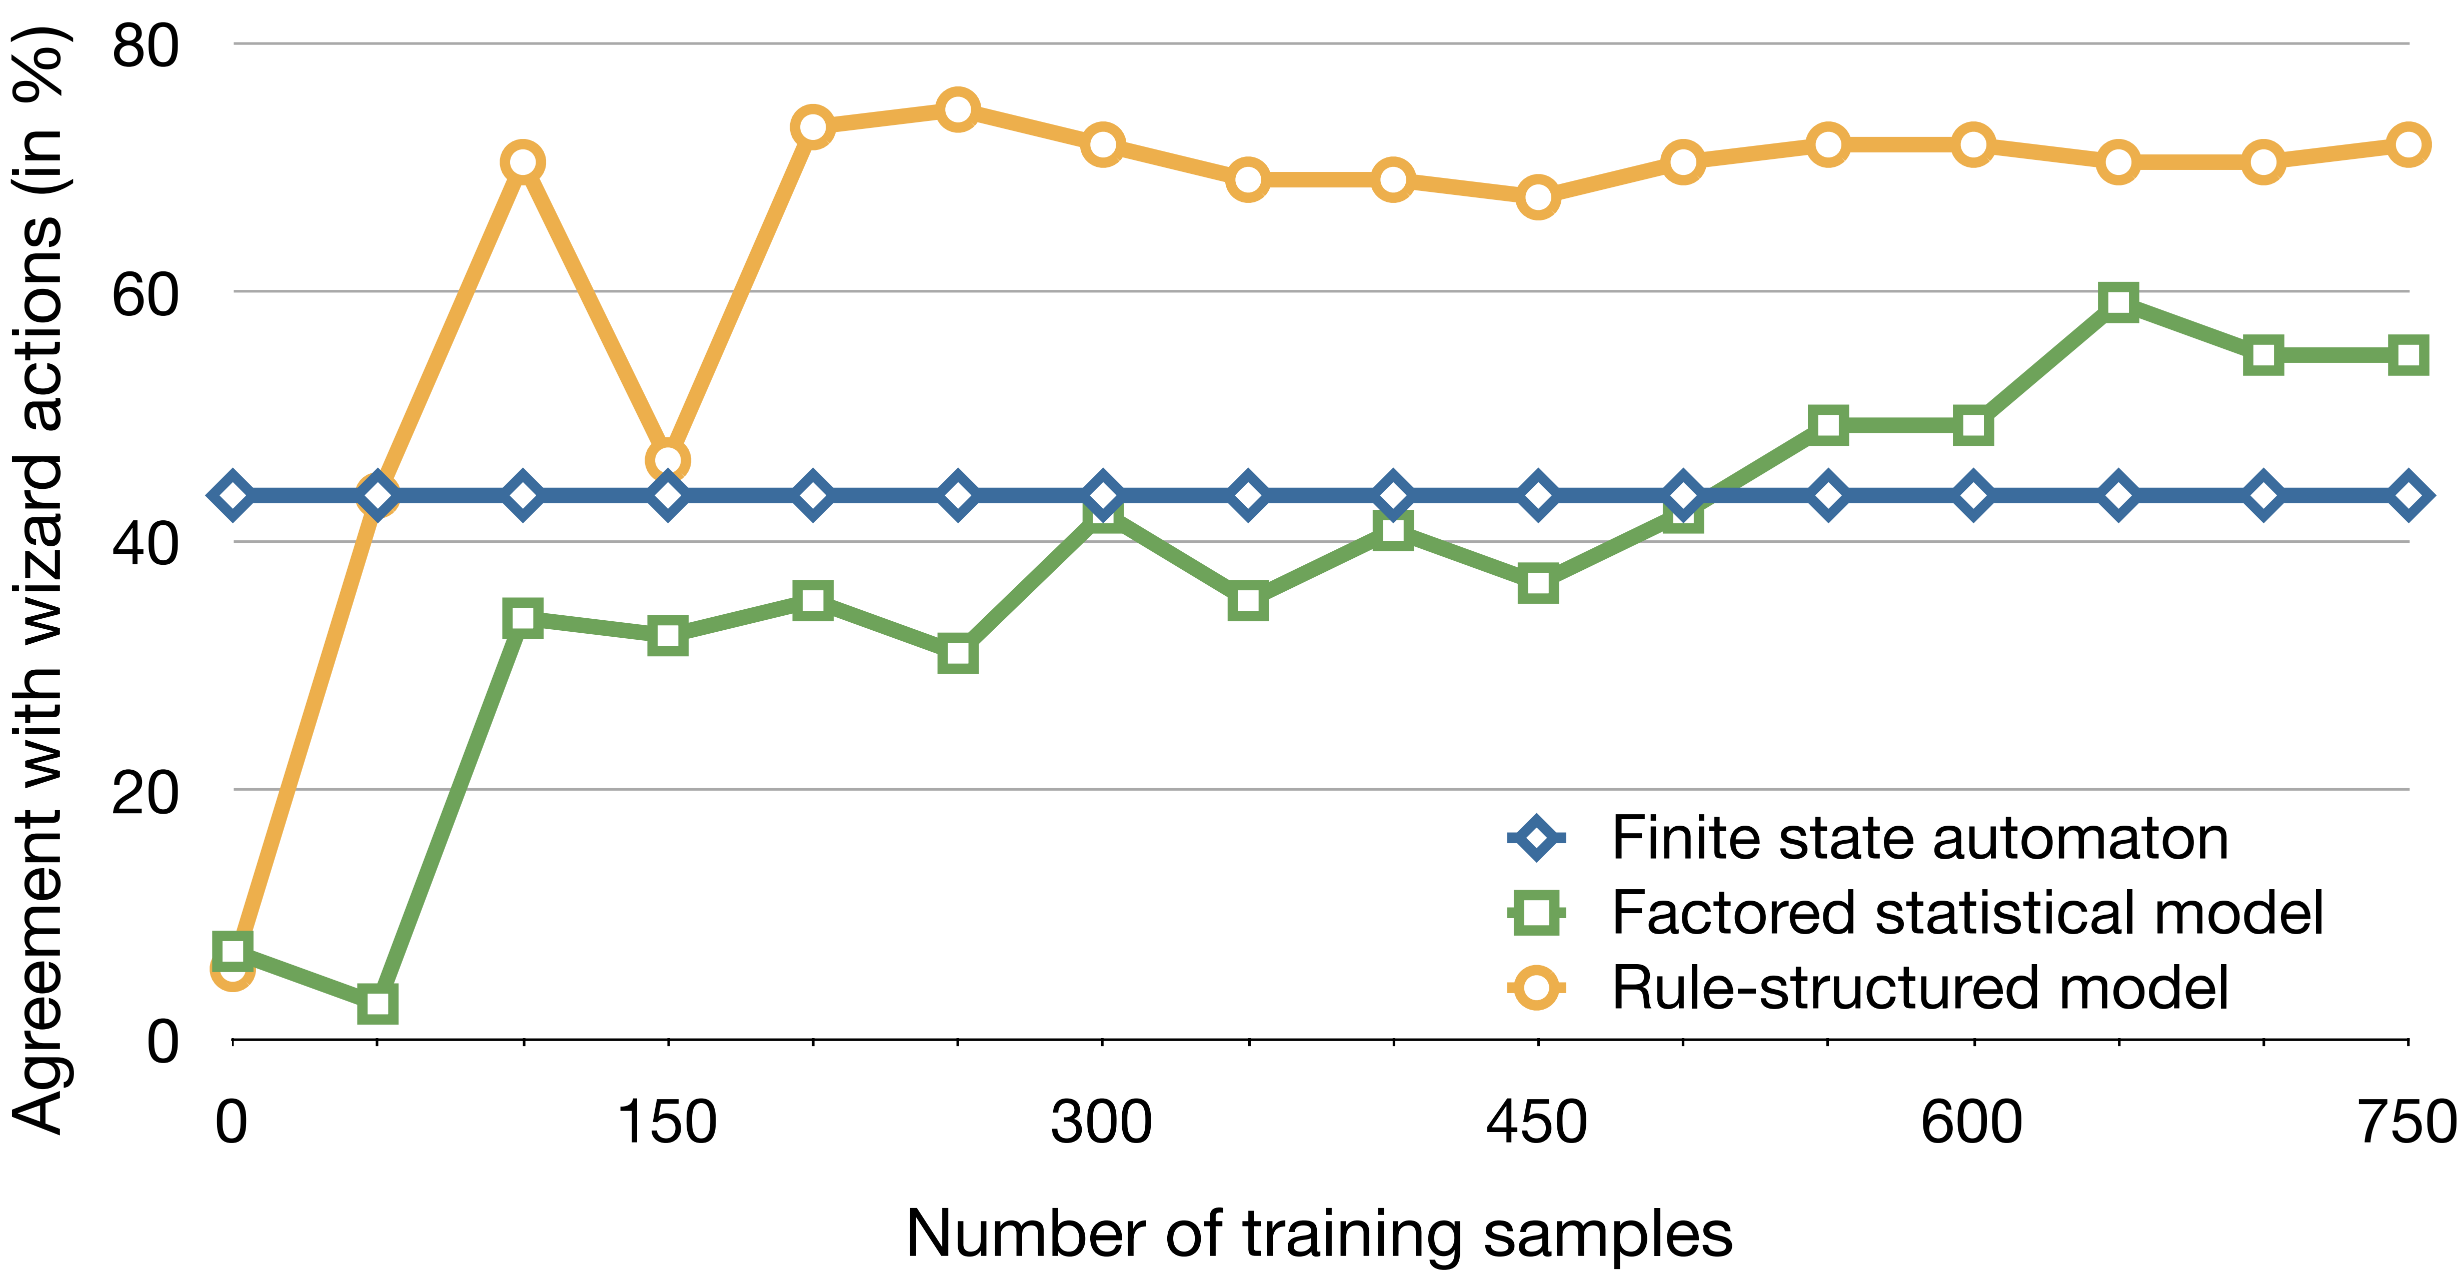

User trials

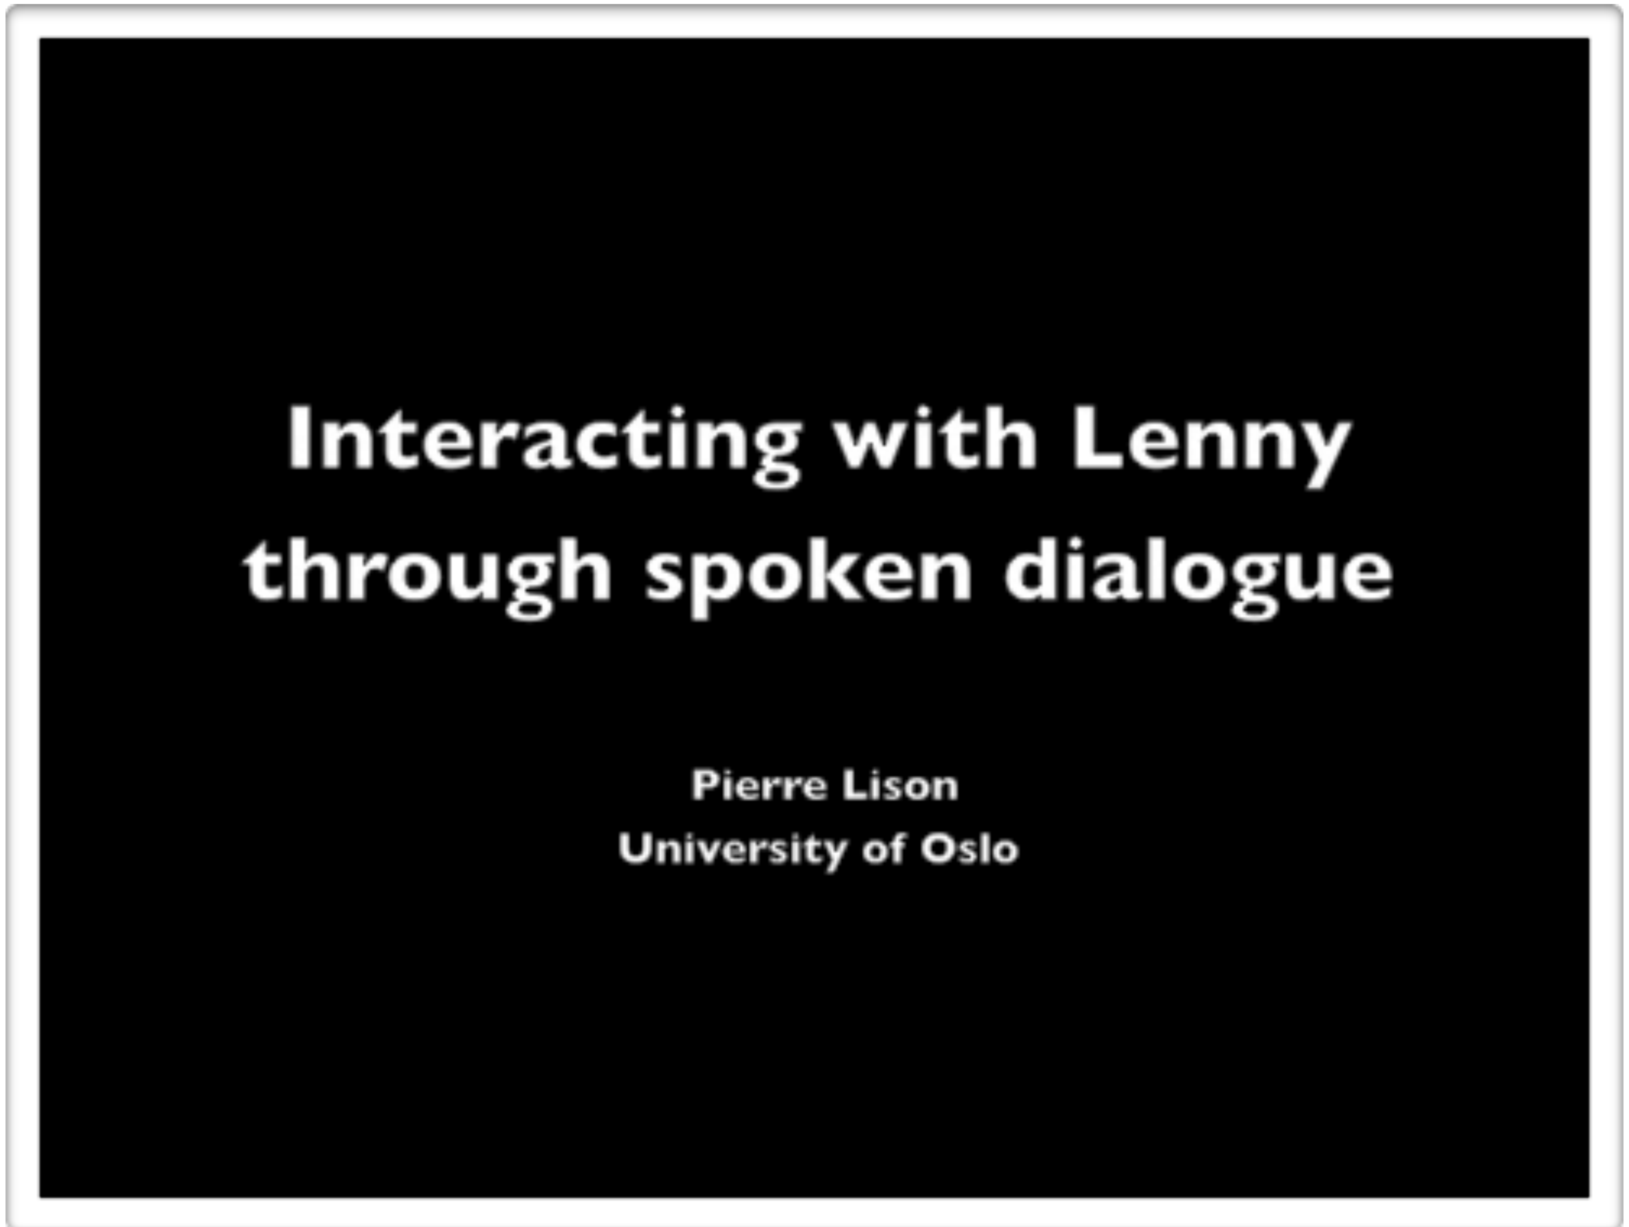

- 37 participants (16 M / 21 F)
- Average age : 30.6
- Average duration: 5:06 mins
- All captured on videos

User trials

- Each participant in the trial repeated the task **three** times
  - One interaction for each modelling approach (in randomised order)
- Evaluation metrics:
  - **Objective metrics**: list of 9 measures extracted from the interaction logs
  - **Subjective metrics** : survey of 6 questions filled by the participants after each interaction

Empirical results

	Metrics	Finite-state automaton	Factored statistical model	Rule-structured model
Objective	Average number of repetition requests	18.68	12.24	**0***
	Average number of confirmation requests	9.16	10.32	**5.78***
	Average number of repeated instructions	3.73	7.97	**2.78**
	Average number of user rejections	**2.16**	2.59	2.59
	Average number of physical movements	**26.68**	29.89	27.08
	Average number of turns between moves	3.63	3.1	**2.54***
	Average number of user turns	78.95	77.3	**69.14**
	Average number of system turns	57.27	54.59	**35.11***
	Average duration (in minutes)	6:18	7:13	**5:24***
Subjective	*"Did you feel that…*			
	… the robot correctly understood what you said?"	3.32	2.92	**3.68**
	… the robot reacted appropriately to your instructions?"	3.70	3.32	**3.86**
	… the robot asked you to repeat/confirm your instructions?"	2.16	2.19	**3.3***
	… the robot sometimes ignored when you were speaking?"	3.24	2.76	**3.43**
	… the robot thought you were talking when you were not?"	3.43	3.14	**4.41***
	… the interaction flowed in a pleasant and natural manner?"	2.97	2.46	**3.32**

Scale from 1 (worse) to 5 (best)

Conclusion

- Development of a new modelling framework for dialogue management, based on **probabilistic rules**
  - *Hybrid* approach at the crossroads between logical and statistical methods
  - Rule parameters can be learned from data
- Experimental studies demonstrate the benefits of the approach
- Concrete implementation in the OpenDial software toolkit