UiO : University of Oslo

Dialogue Management with Probabilistic Rules

Pierre Lison
Language Technology Group
University of Oslo

Göteborg, 20th May 2014

What is dialogue management?

Component of a *spoken dialogue system*, in charge of the following tasks:

- maintain a representation of the *dialogue state*
- Select the *best action* to execute

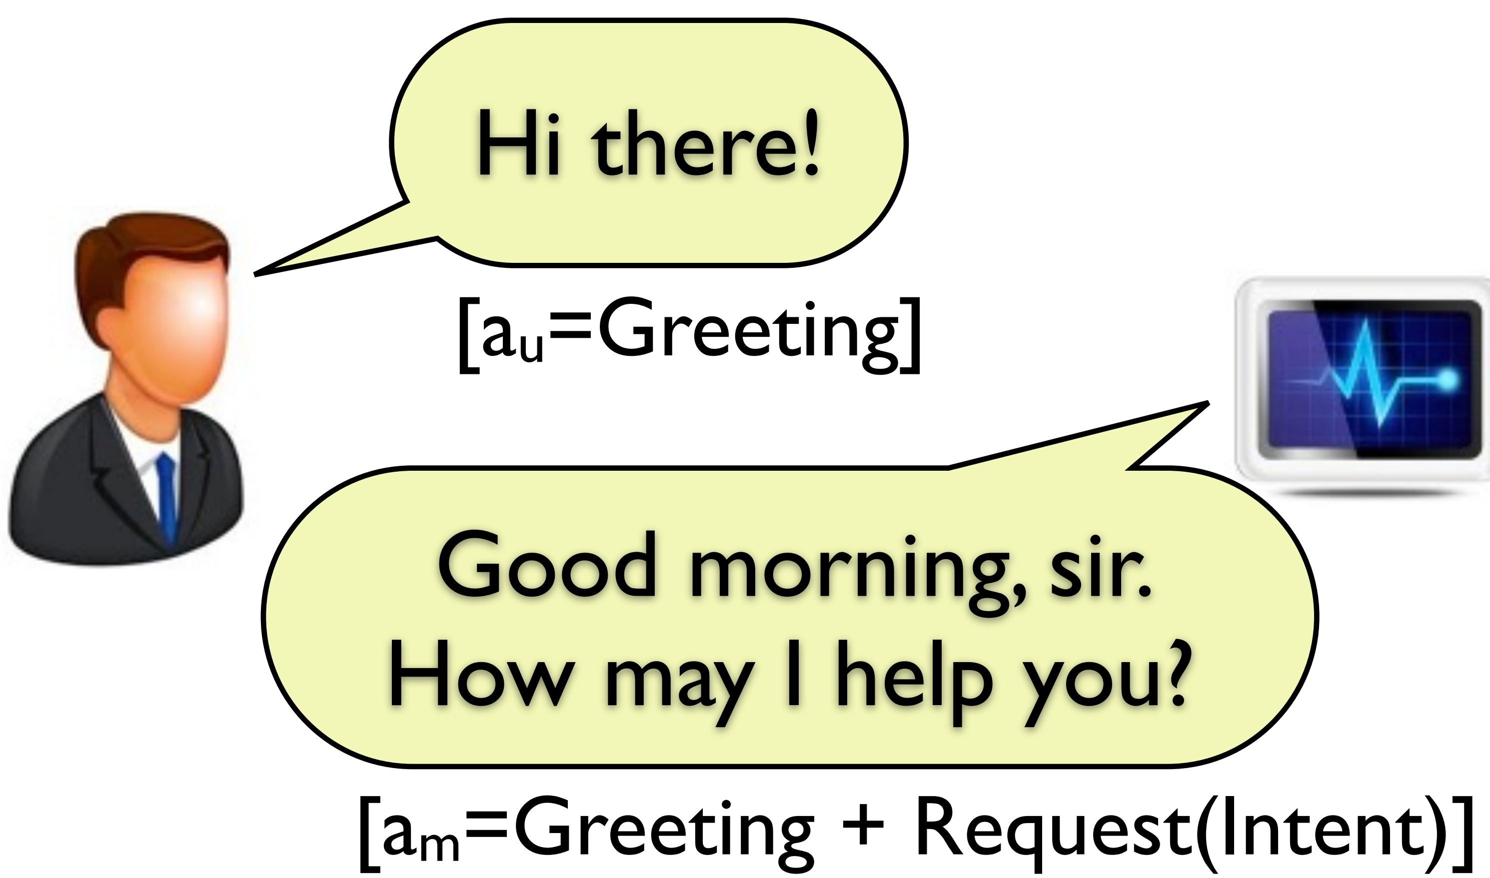

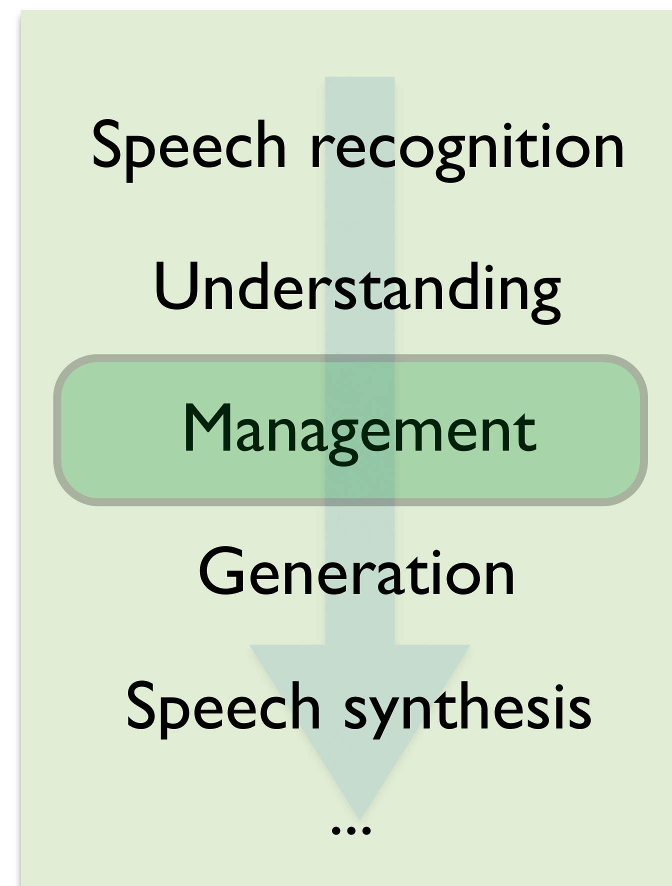

Challenges

Spoken dialogue is ...

Complex

- Context is essential to understand many utterances
- 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*

Outline

1. Probabilistic rules
2. Dialogue modelling
3. Parameter estimation
4. Experiments
5. Conclusion

Modelling approach

- **Goal**: combine the best of *probabilistic* and *logical* approaches to dialogue management
- Dialogue state represented as a *Bayesian network*
  - Interconnected set of variables, where each captures a relevant aspect of the interaction
  - regularly updated with new observations & used to derive high-utility actions
- **Key idea**: Encode domain models with a rich, structured modelling language

Probabilistic rules

- A probabilistic rule specifies a particular *relation* between state variables
  - Mapping between *conditions* and (probabilistic) *effects*
  - Can use logical operators and quantifiers
  - Structured as an *if…then…else* construction:

if (condition$_1$ holds) **then**
 …
else if (condition$_2$ holds) **then**
 …
else
 …

Types of rules

	Probability rules	**Utility rules**
What they encode:	*Conditional probability distributions* between state variables	*Utility distributions* for system actions given state variables
General skeleton:	**if** (condition$_1$) **then** P(effect$_1$)= θ_1, P(effect$_2$)= θ_2, ... **else if** (condition$_2$) **then** P(effect$_3$) = θ_3, ...	**if** (condition$_1$) **then** U(action$_1$)= θ_1, U(action$_2$)= θ_2, ... **else if** (condition$_2$) **then** U(action$_3$) = θ_3,... ...

Examples of probabilistic rules

Example (rule r_1):

$\forall x$, **if** (a_m = AskRepeat $\wedge$ a_u = x) **then**
  $P(a_u' = x) = 0.9$

"If the system asks the user to repeat his last dialogue act *x*, the user is predicted to comply and repeat *x* with probability 0.9"

Example (rule r_2):

$\forall x$, **if** (a_u=Request(PickUp(x)) $\wedge$ $x \in$ perceived) **then**
  $U(a_m'$=Do(PickUp(x))) = +5

"If the user asks the system to pick up a given object *x* and *x* is perceived by the system, then the utility of picking up the object is 5"

Rule instantiation

Probabilistic rules are **high-level templates** for a (directed) graphical model

Example (rule r_1):

$\forall x$, **if** (a_m = AskRepeat $\wedge$ $a_u = x$) **then**
$P(a_u' = x) = 0.9$

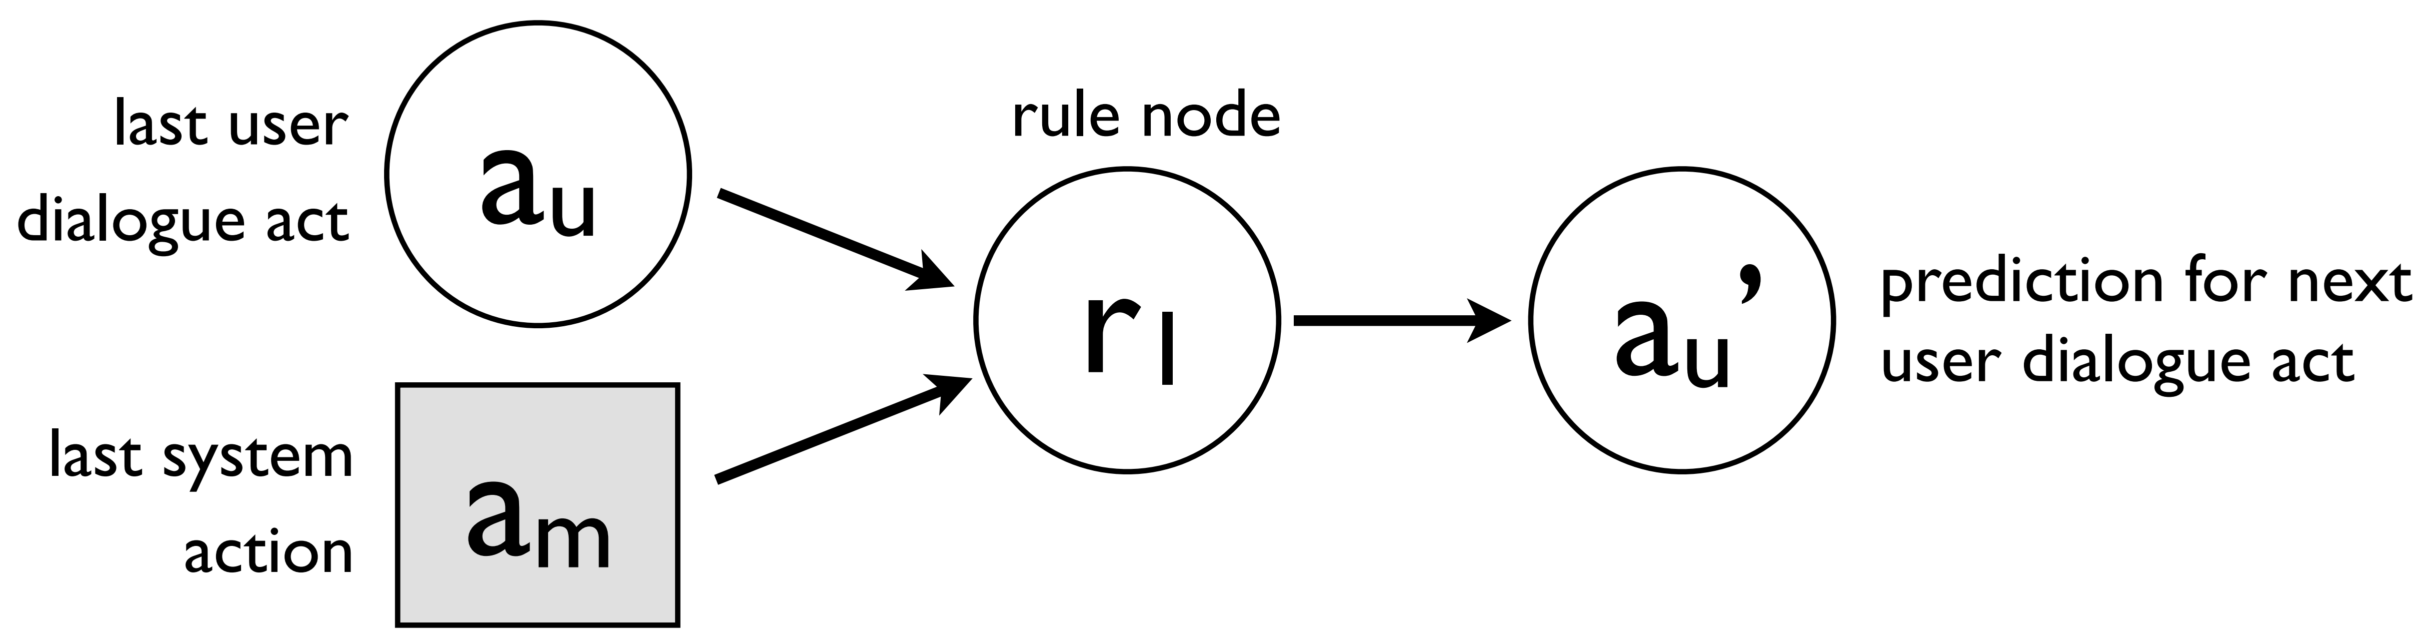

Rule instantiation

Probabilistic rules are **high-level templates** for a (directed) graphical model

Example (rule r_2):

$\forall x$, **if** (a_u=Request(PickUp(x)) $\wedge$ $x \in$ perceived) **then**

U(a_m'=Do(PickUp(x))) = +5

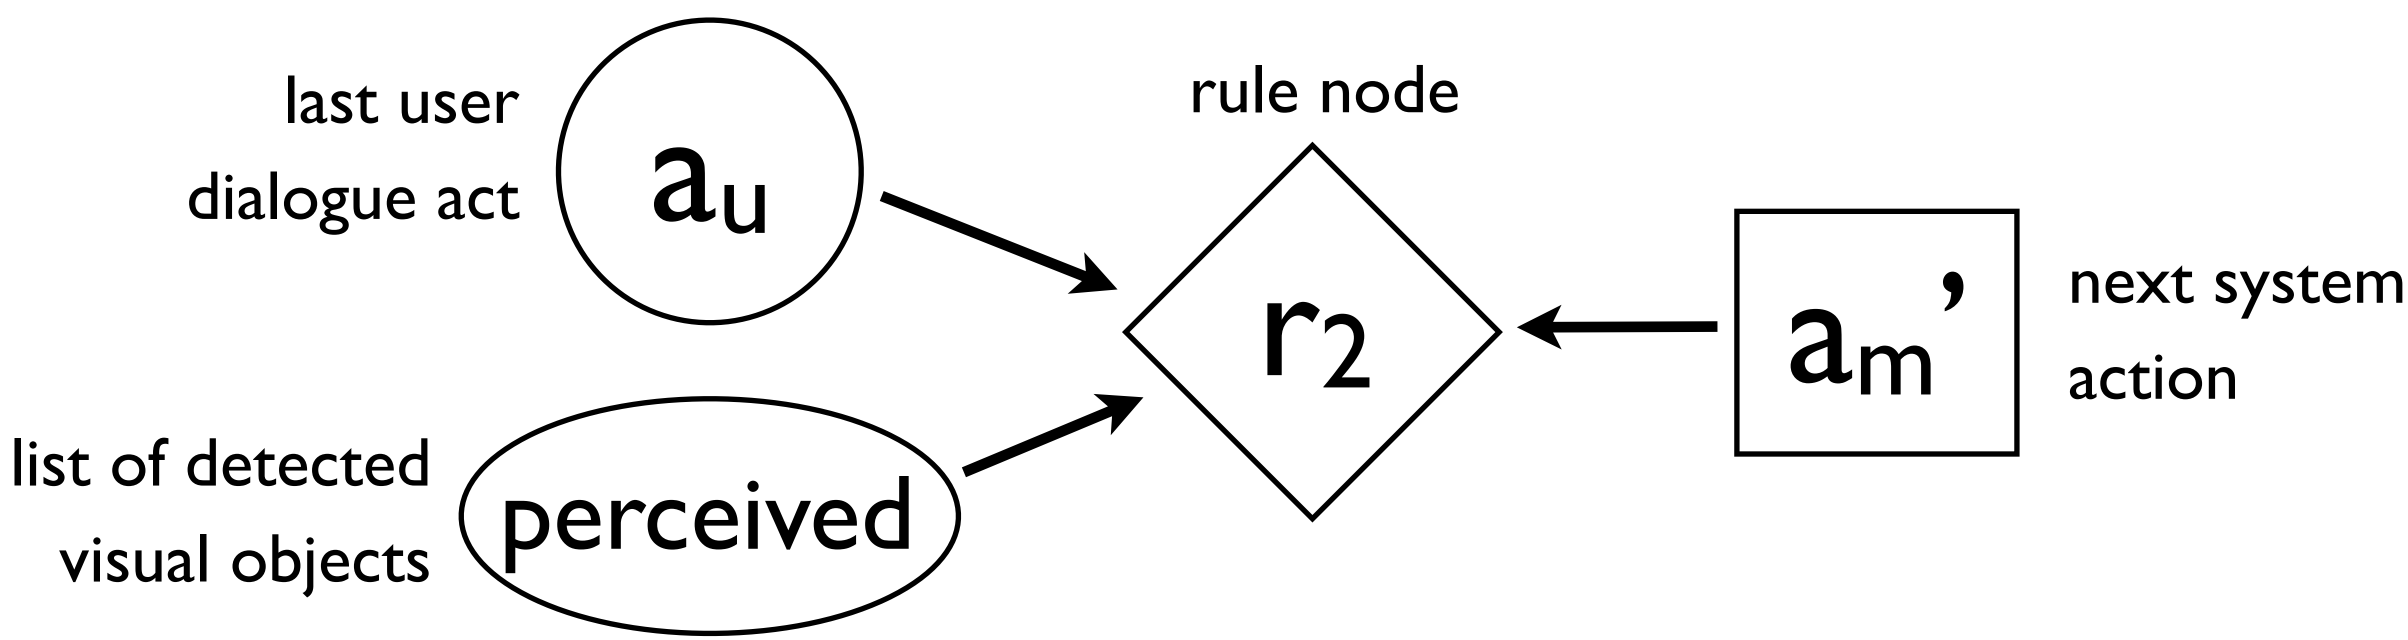

Probabilistic rules for dialogue

Claim: Probabilistic rules are well suited to structure the probability and utility models employed in dialogue management

Ability to learn domain parameters with limited amounts of data

Ability to capture expert domain knowledge in human-readable form

Processing workflow

- *Blackboard architecture* centered around the dialogue state
- Rule-structured models and other modules can read & write to the dialogue state
- Implementation in the OpenDial toolkit [http://opendial.googlecode.com]

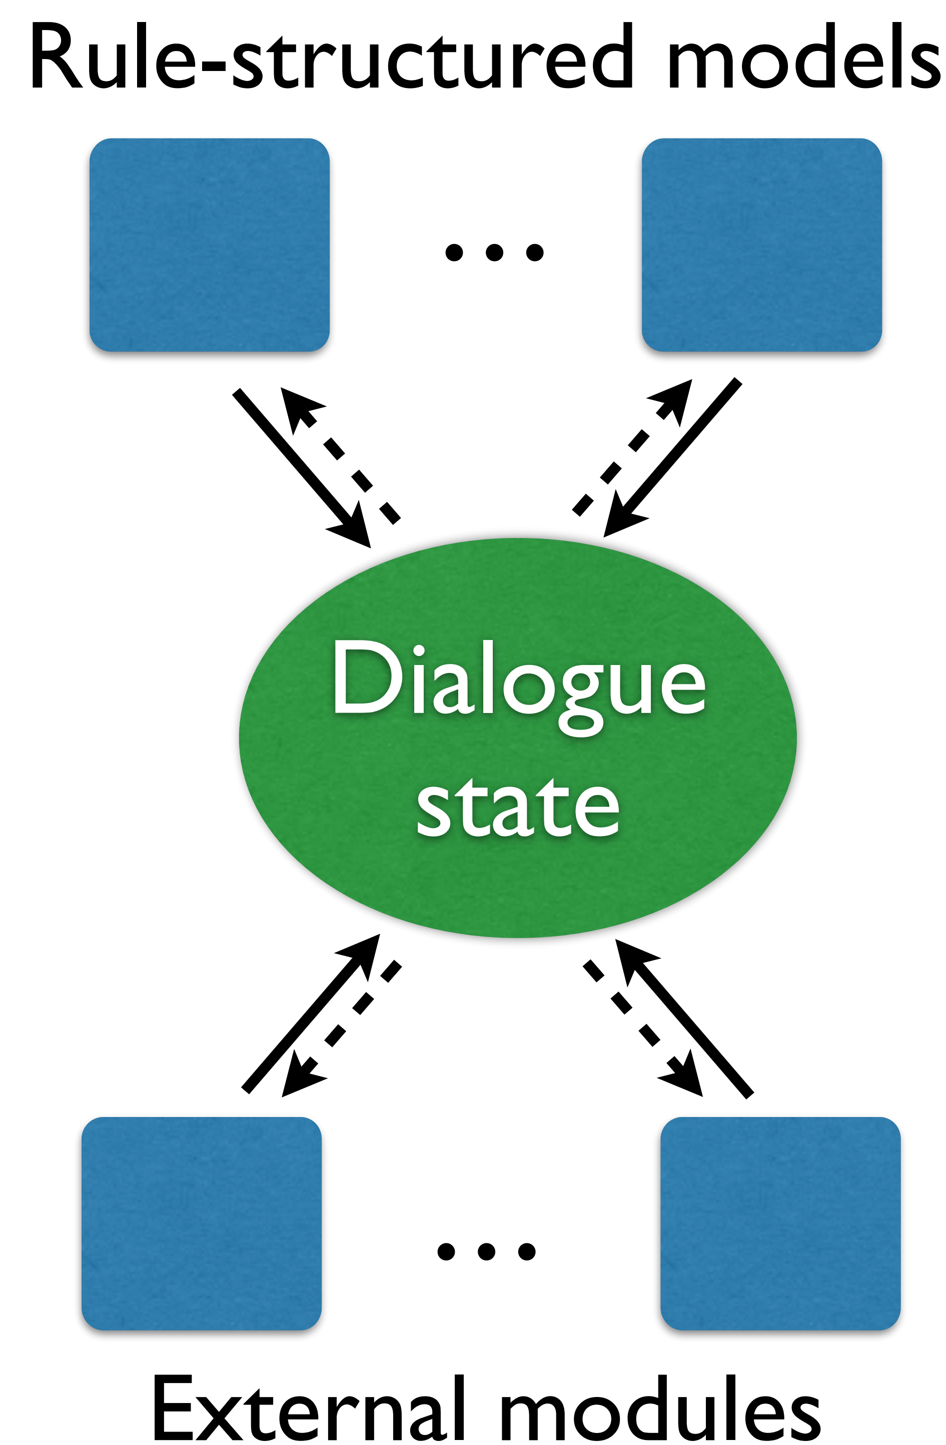

Parameter estimation

- 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 $\mathcal{D}$

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

Instantiation of parameter nodes

Example (rule r_1):

$\forall x$, **if** (a_m = AskRepeat $\wedge$ $a_u = x$) **then**
$P(a_u^p = x) = \theta$

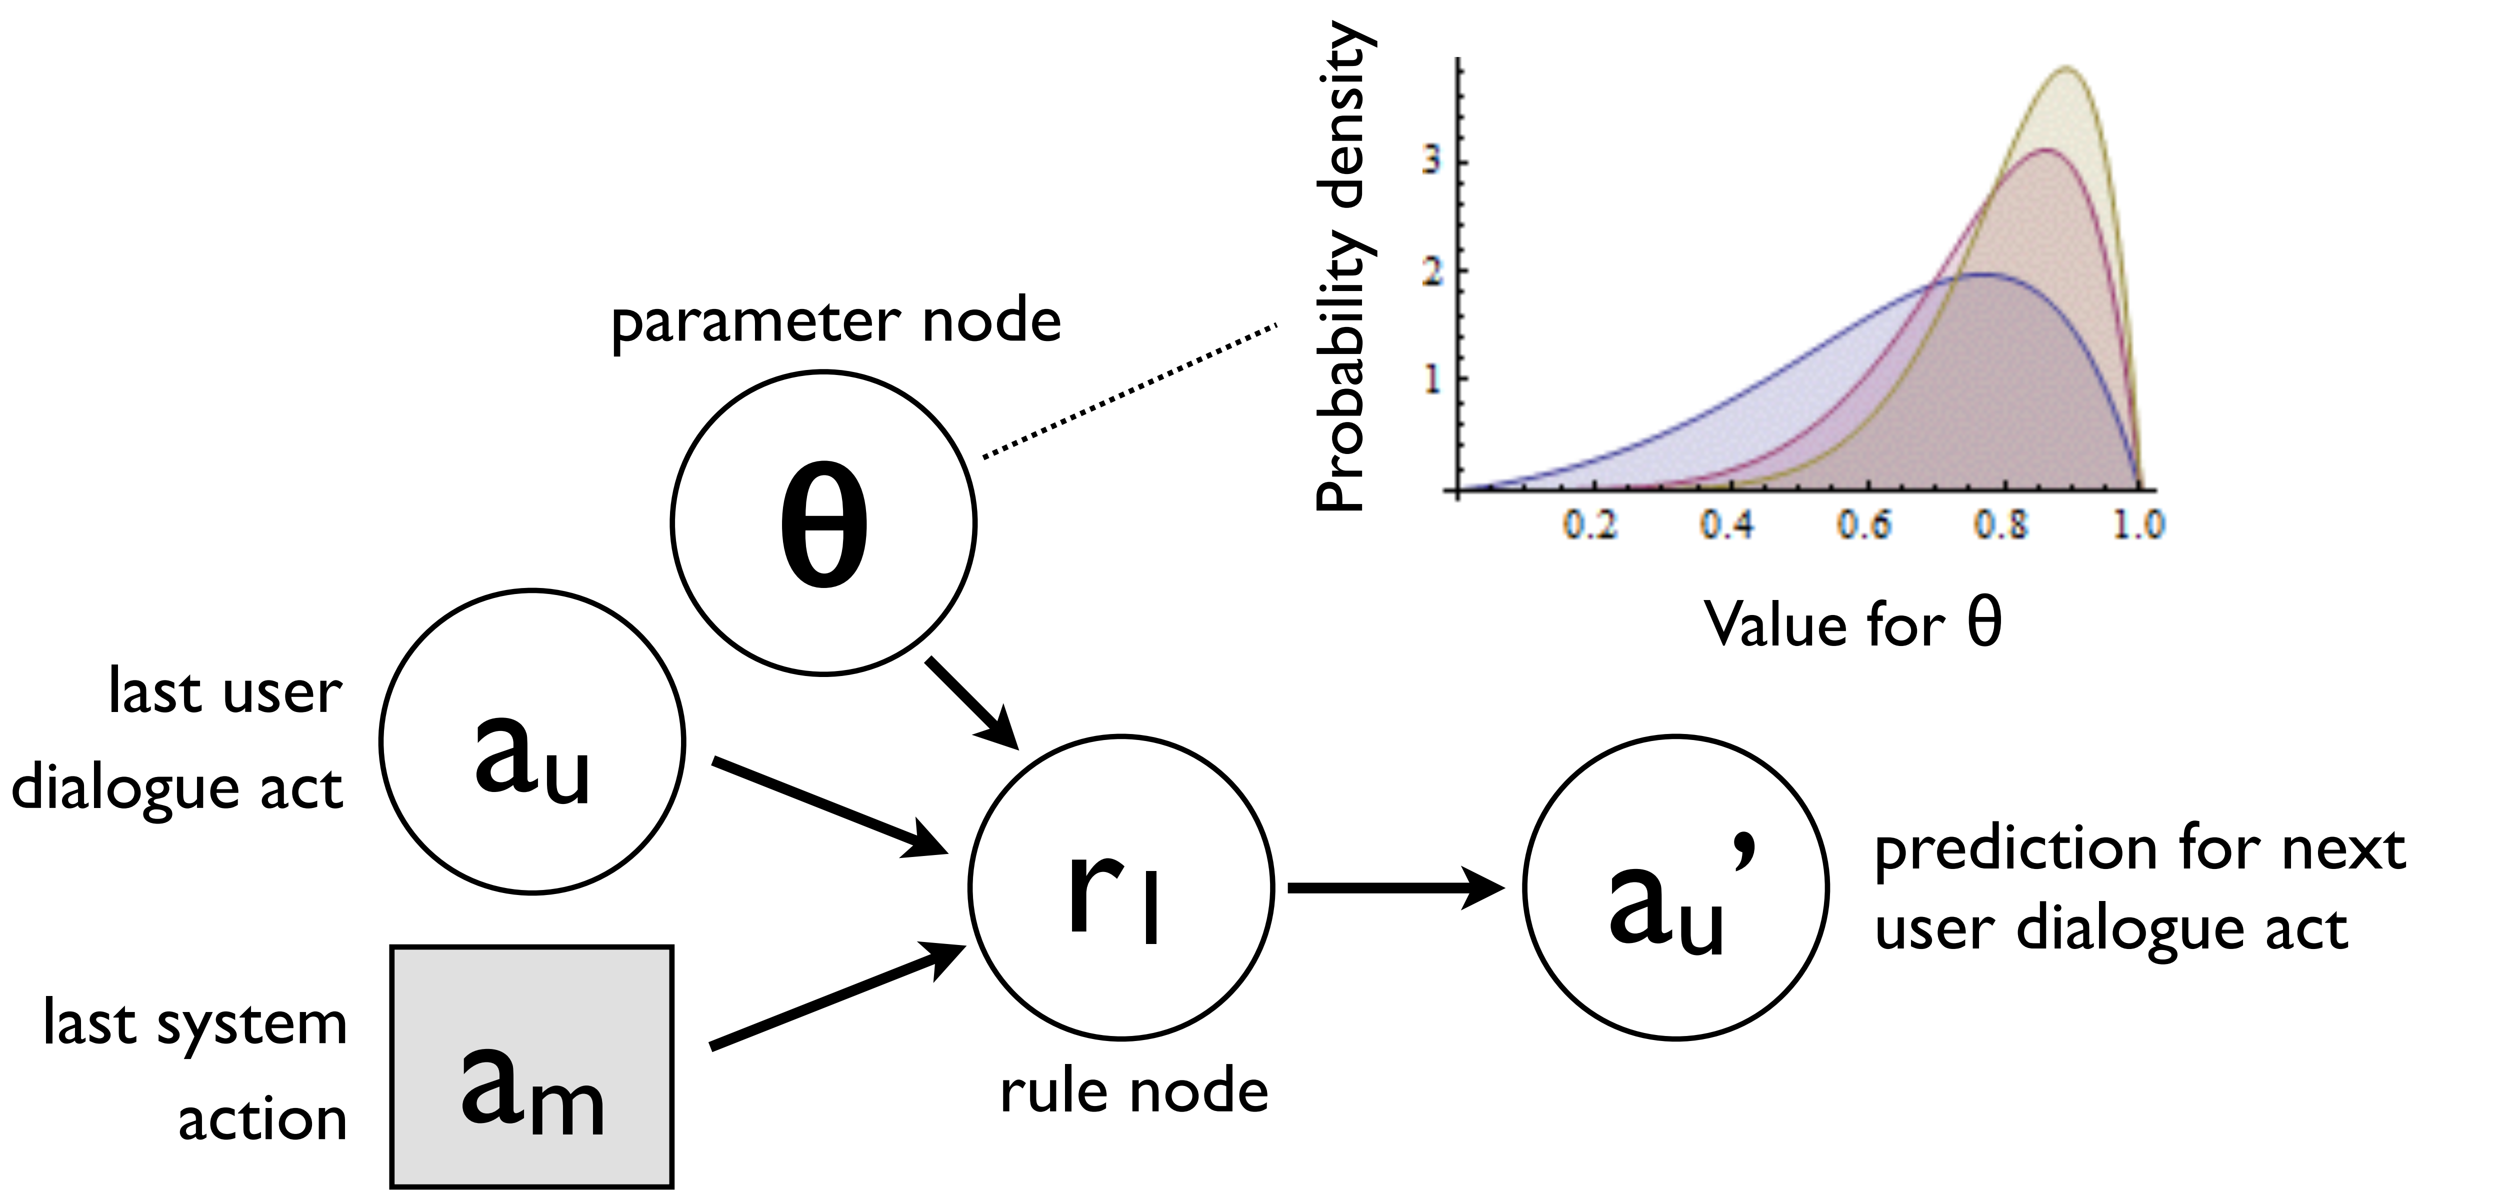

Learning paradigms

- 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

Live example with OpenDial

Learning the probability θ that the user will repeat his last dialogue act when asked by the dialogue system to do so

- Initial prior distribution: *Beta(2,1)*
- The parameter distribution is automatically updated after each observation
- The distribution is gradually narrowed down to a region of the parameter space

Experiments

- Several experiments conducted to assess the viability of the modelling framework:
  - Analysis of learning performance on a small Wizard-of-Oz data set
  - Analysis of learning performance with a user simulator
  - Empirical evaluation of dialogue quality with a user trial with 37 participants

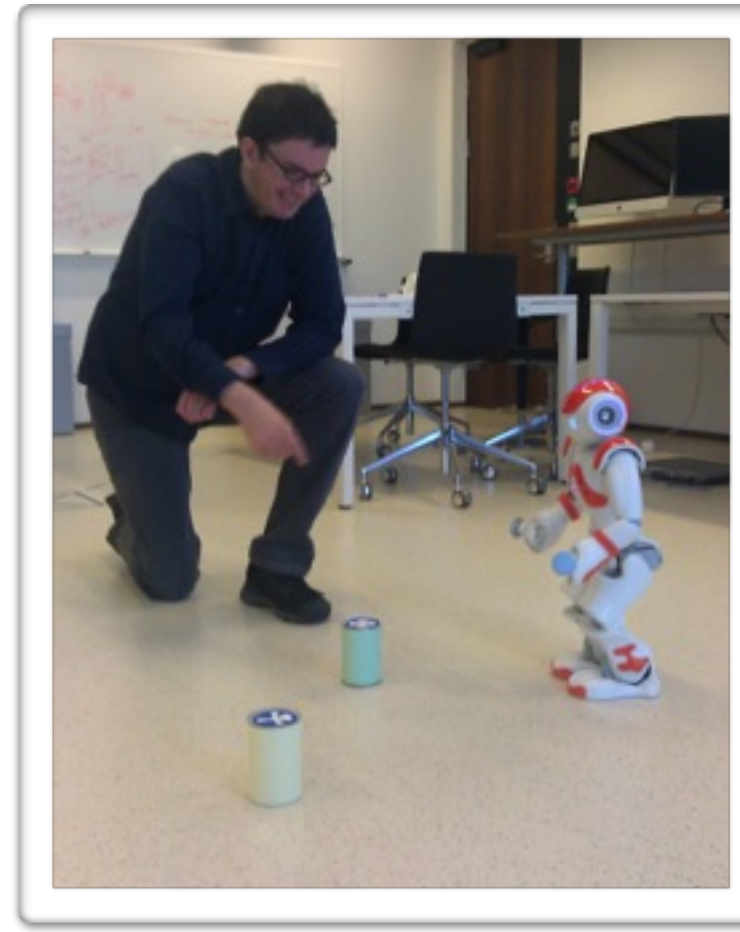

[P. Lison. Probabilistic Dialogue Models with Prior Domain Knowledge (SIGDIAL 2012)]
[P. Lison. Model-based Bayesian Reinforcement Learning for Dialogue Management (Interspeech 2013)]

User evaluation

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

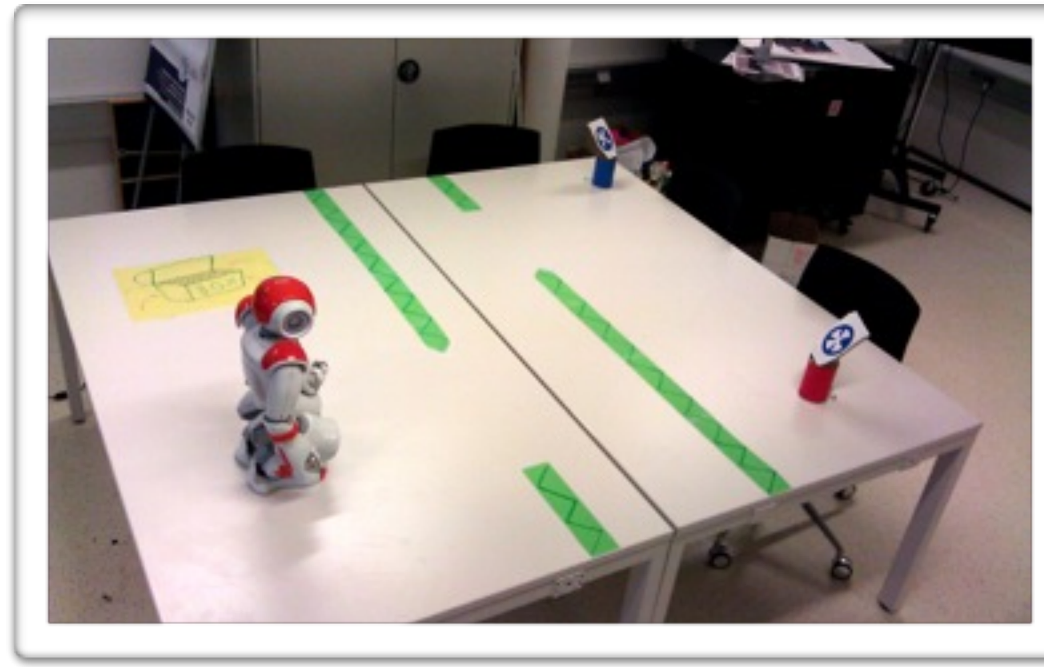

- Comparison of three modelling approaches:

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: 335×10^6 (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)

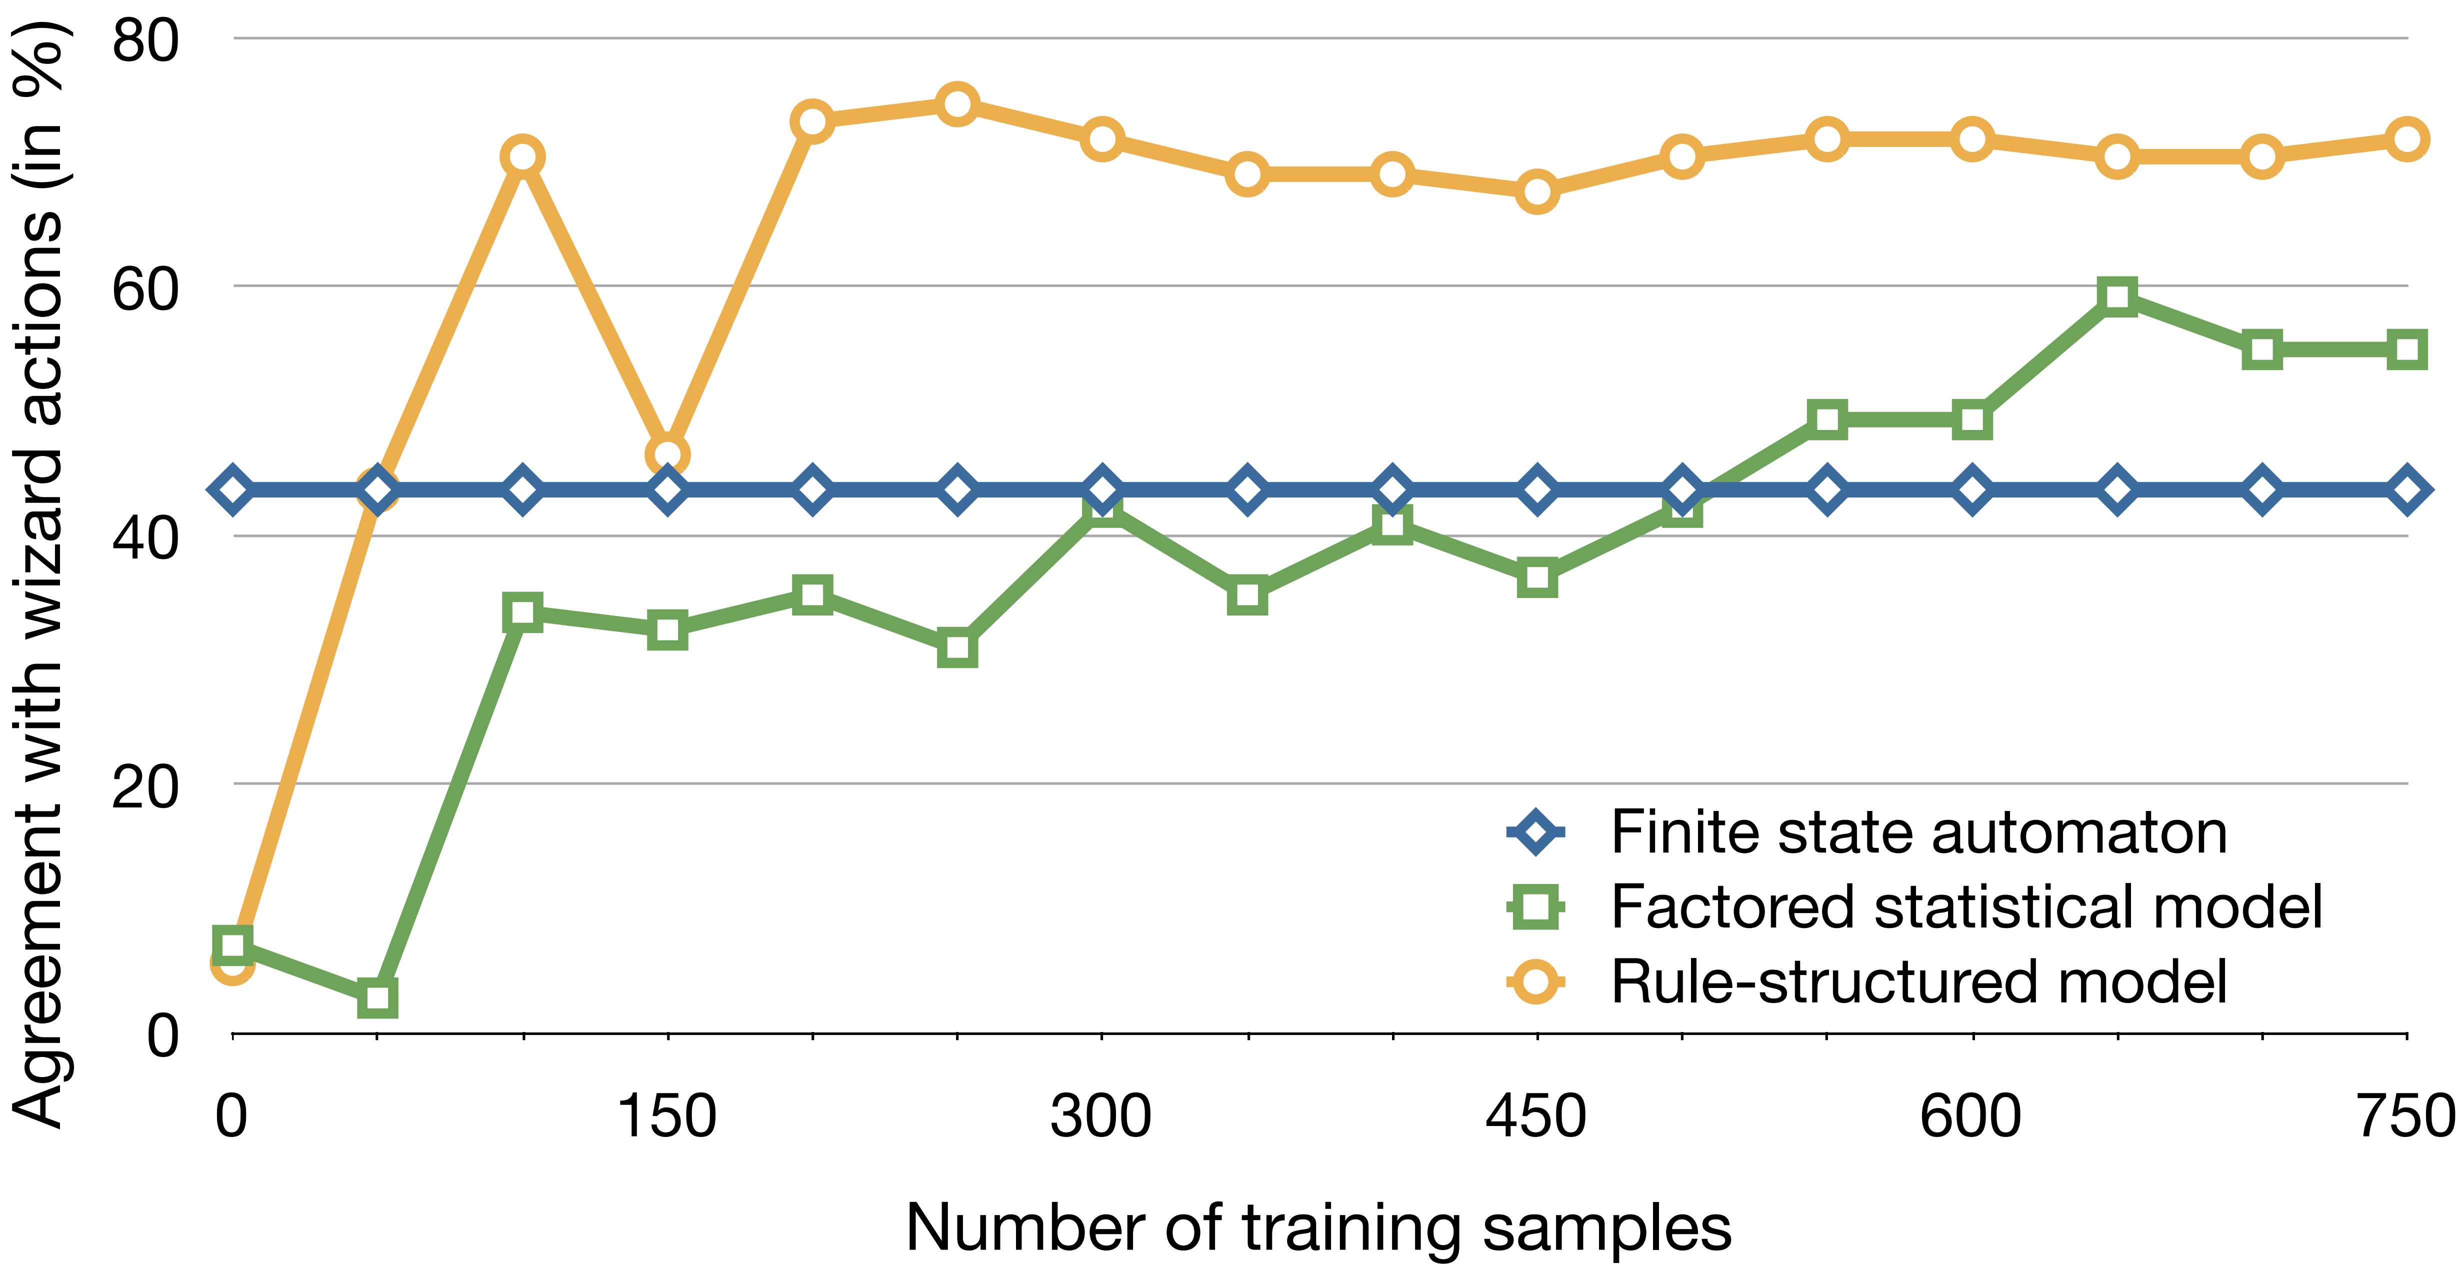

User trials

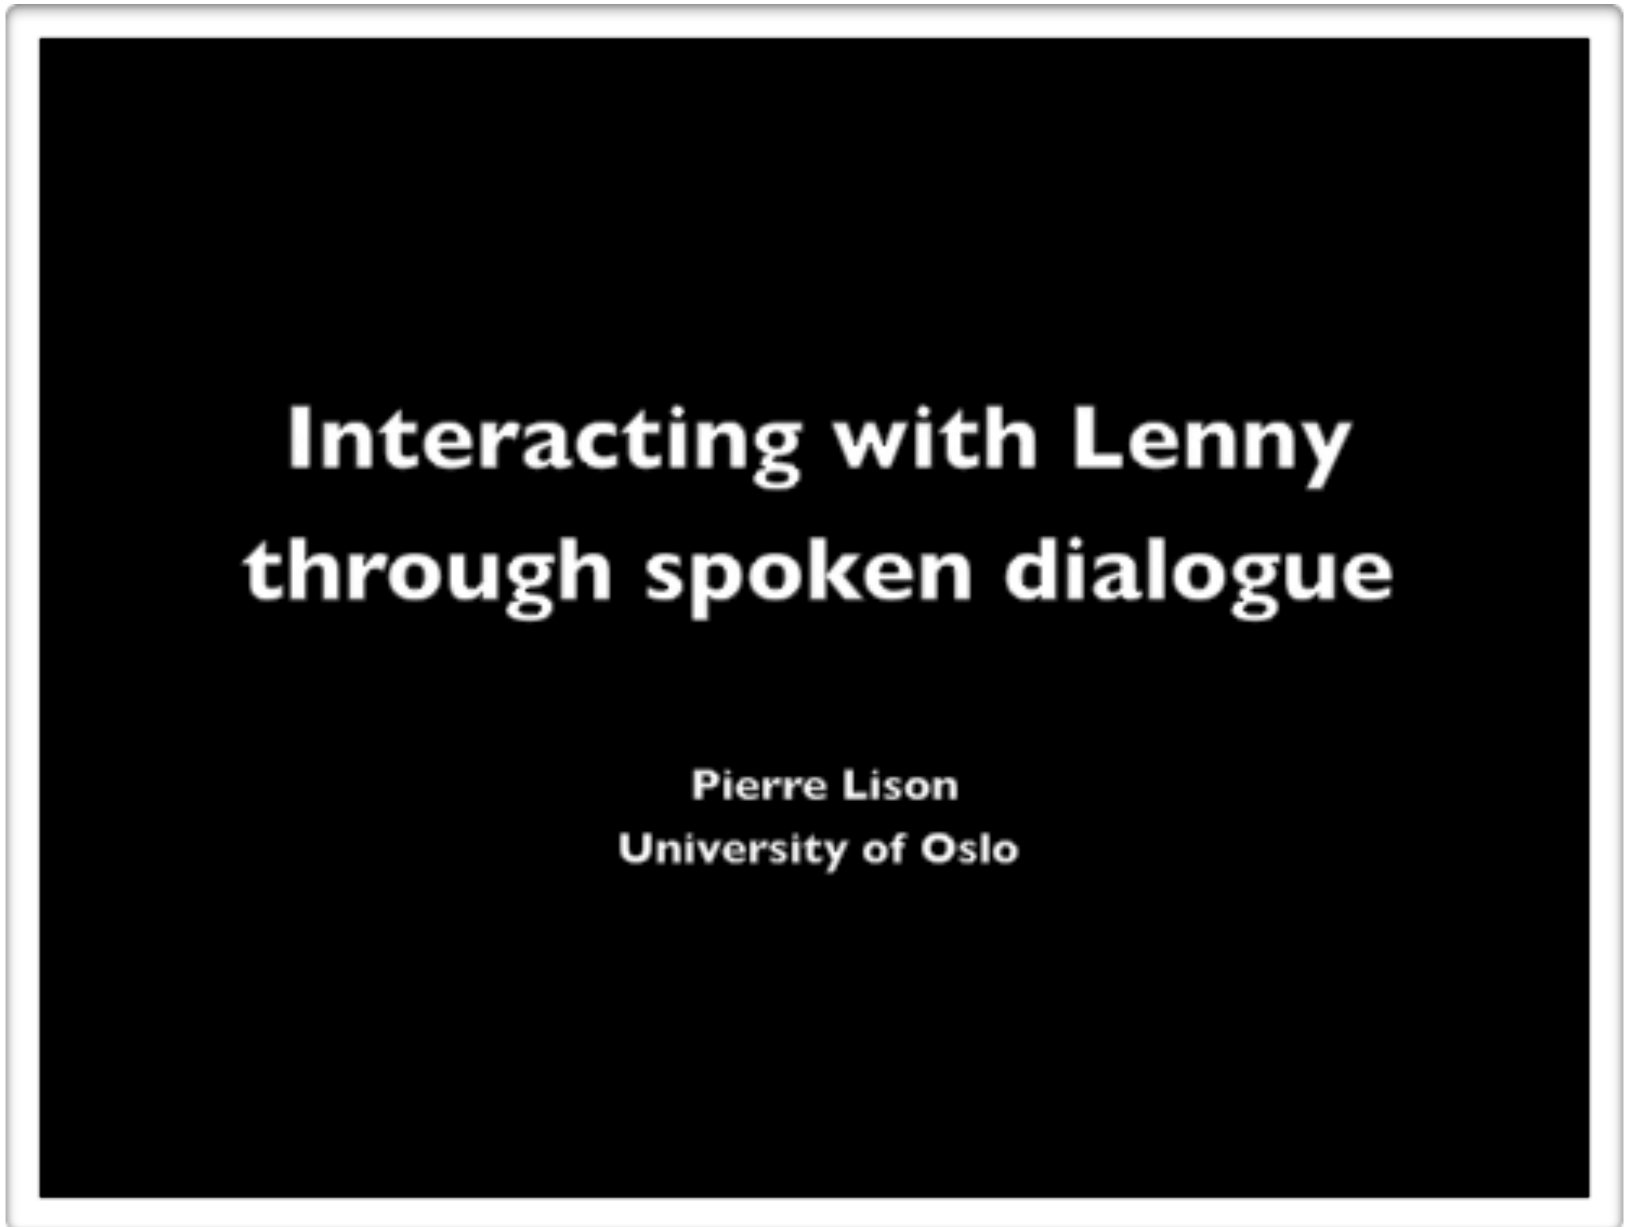

- 37 participants (16 M / 21 F)
- Average age : 30.6
- Average duration: 5:06 mins
- All captured on videos

User trials

- Each of the 37 participants 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