



Towards Dialogue Management in Relational Domains

Pierre Lison,
Language Technology Group (LTG)
Department of Informatics

Workshop on Action, Perception & Language (APL)

October 25 2012



Introduction

- *Slot-filling applications* are still often considered as «prototypical» domains for dialogue systems
- Their representation of the dialogue state is based on particular modelling assumptions:

Task	One single task: filling the predefined slots
Context	None (or very limited): no external context to capture
User model	None (or very limited): different user for each interaction

Introduction

- But these modelling assumptions do not hold for many other domains - especially *situated domains*
- Ex: human-robot interaction, cognitive assistants & companions, tutoring systems, etc.




Task	One single task: filling the predefined slots	<i>Varying number of interconnected tasks</i>
Context	None (or very limited): no external context to capture	<i>Rich, dynamic, often situated environment</i>
User model	None (or very limited): different user for each interaction	<i>longer interactions → complex user modelling</i>



Introduction

- These situated domains have a rich internal structure
- This structure is often best described in terms of *entities* and *relations* between entities:
 - Physical objects spatially connected in a visual scene
 - Indoor environments with places in which to navigate
 - Stacks of (interconnected) tasks to complete



Dialogue state expressed as a
relational structure



The problem

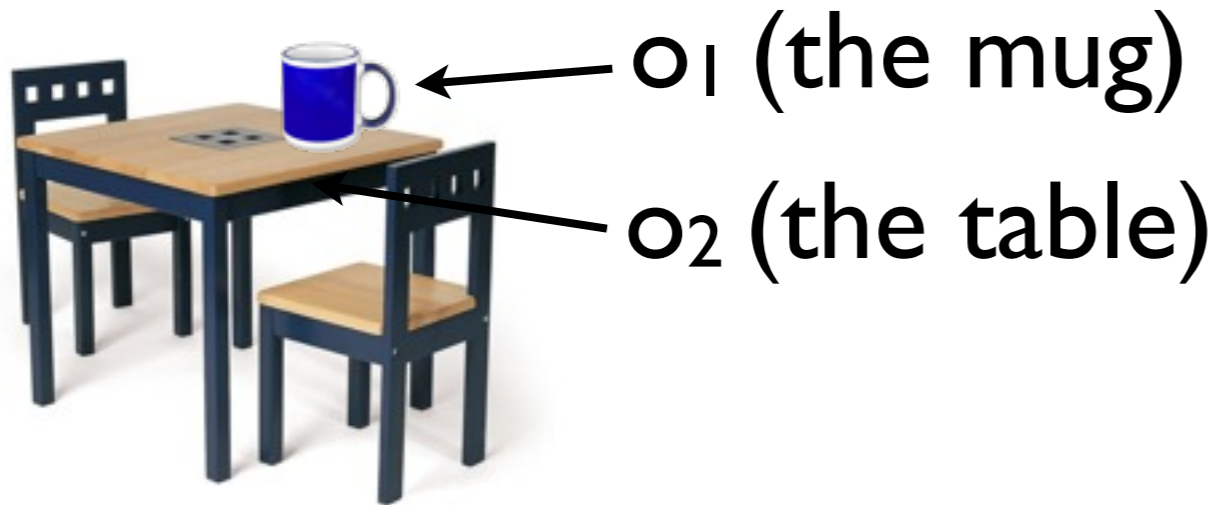
- **Starting point:** we wish to encode the dialogue state as a *relational structure*
- **Problem 1:** how to represent this dialogue state in practice, without giving up the probabilistic modelling?
- **Problem 2:** how do we build *dialogue models* (for interpretation & action selection) that can operate on such state representation?



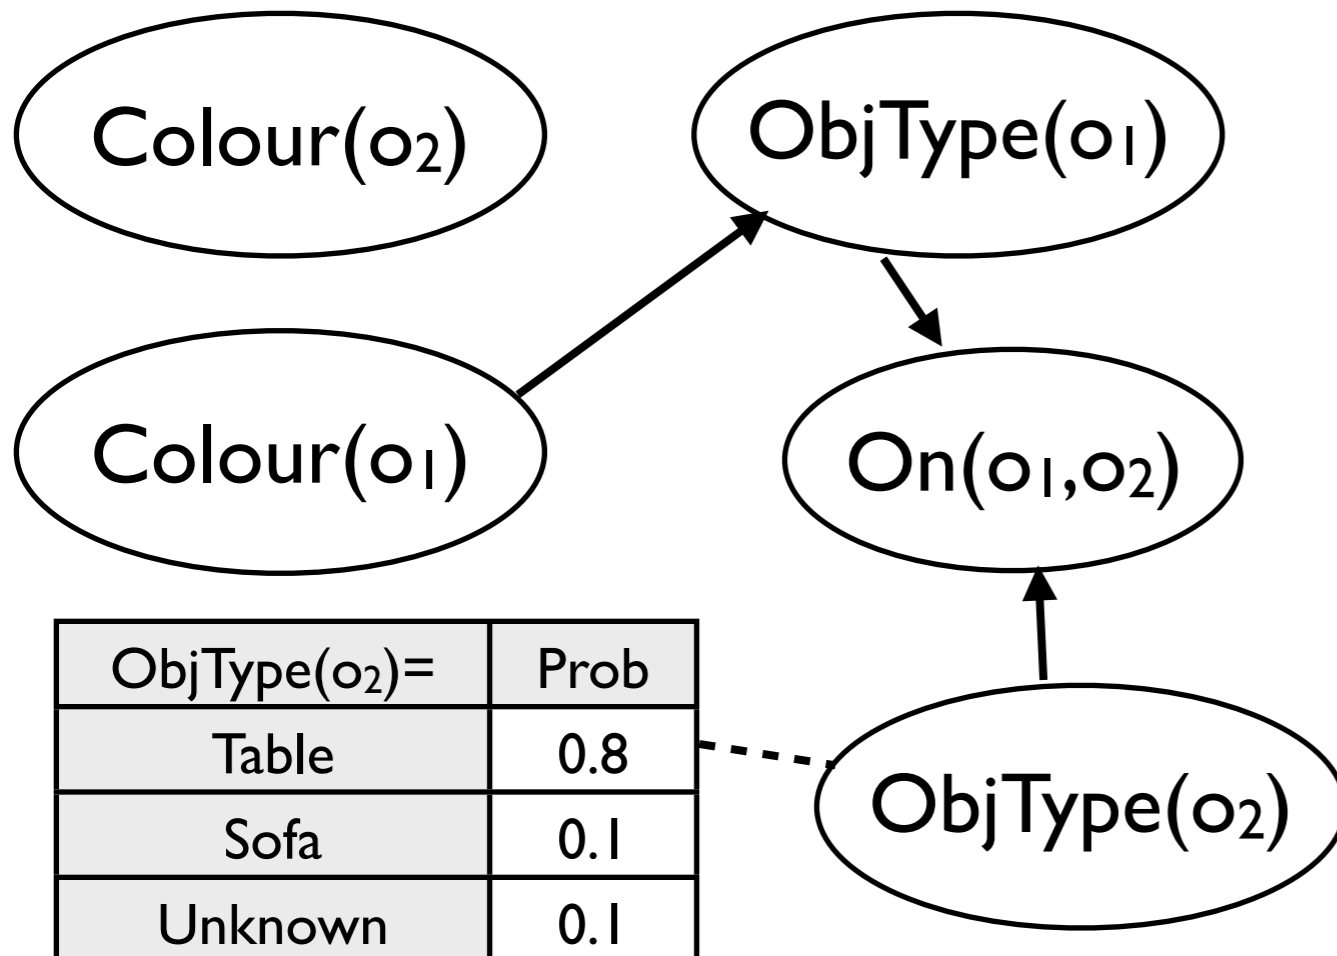
Problem 1: dialogue state

- **Problem 1:** We need a representation of the dialogue state that is able to:
 - capture the domain's relational structure of the domain
 - account for the pervasive uncertainty in spoken dialogue, especially in situated domains
- We encode our dialogue state as a **Bayesian Network** (i.e. a *directed graphical model*)
- The variables of this network are *grounded predicates and functions*

Problem 1: dialogue state



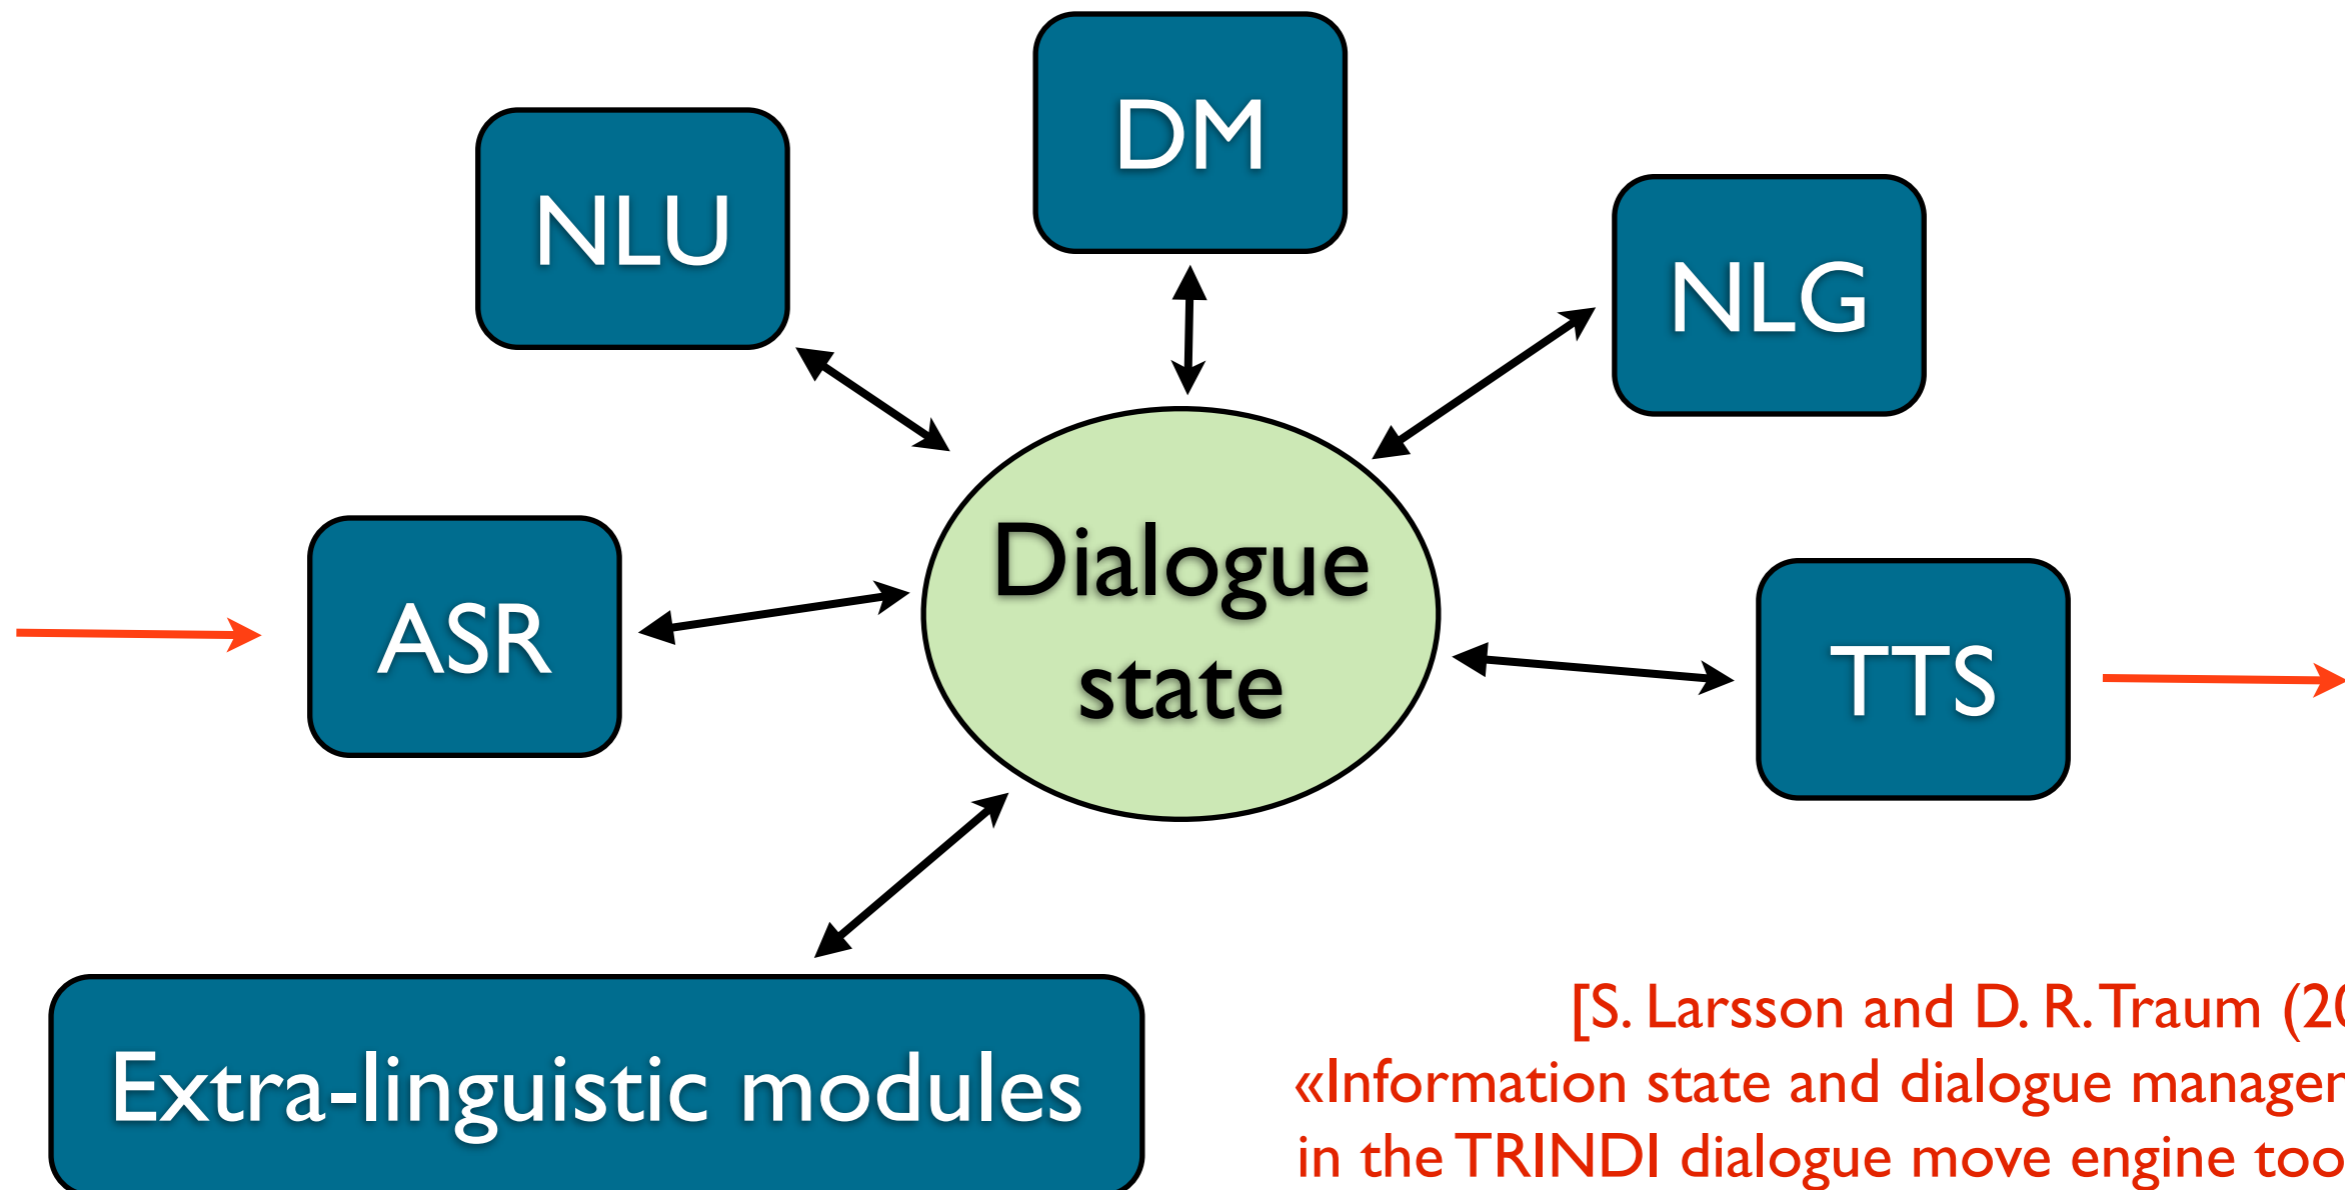
Assume we want to encode in our dialogue state the type & colour of these two objects, as well as their relative spatial position



- The variable labels are ground predicates or functions
- Each variable is associated with a *probability distribution* defining its possible values
- The variable values might depend on each other (*conditional dependencies*)

Problem 1: dialogue state

- *Information-state* architecture: the state acts as a *blackboard* read and written by the various processing components



[S. Larsson and D. R. Traum (2000),
«Information state and dialogue management
in the TRINDI dialogue move engine toolkit»
in *Natural Language Engineering*]



Problem 2: dialogue models

- **Problem 2:** how do we build practical dialogue models defined on the dialogue state representation we just presented?
 - For dialogue interpretation, action selection, etc.
- What we want:
 - *Probabilistic models* of dialogue processing
 - ... that have *parameters* that can be estimated from data
 - ... and take advantage of the domain's *internal structure*

Problem 2: our solution

- **Proposed solution:** encode the dialogue models via expressive *probabilistic rules* with a limited form of *quantification*

Approach grounded in *probabilistic modelling*:

- principled account of uncertainties
- parameters can be estimated from data

... but also employing *expressive rules* that can compactly capture

- high-level generalisations
- prior domain knowledge

... and that can range over arbitrary *sets of entities*



Probabilistic rules

- Probabilistic rules take the form of structured **if...then...else** cases
- Mapping from *conditions* to (probabilistic) *effects*:

```
if (condition1 holds) then  
  P(effect1) =  $\theta_1$ ,   P(effect2) =  $\theta_2$ ,   ...  
else if (condition2 holds) then  
  P(effect3) =  $\theta_3$ ,   ...  
...  
else  
  P(effectn) =  $\theta_n$ ,   ...
```



Probabilistic rules

- *Conditions* are (arbitrarily complex) logical formulae on state variables
- *Effects* are value assignments on state variables
- Effect probabilities are *parameters* that can be estimated from data

Example:

if ($a_m = \text{AskRepeat}$) **then**

$$P(a_u' = a_u) = 0.9$$

$$P(a_u' \neq a_u) = 0.1$$



Utility rules

- The formalism can also describe *utility models*
- In this case, the rule maps each condition to an assignment of *utility values* for particular actions:

if (condition₁ holds) **then**

$Q(\text{actions}_1) = \theta_1, Q(\text{actions}_2) = \theta_2, \dots$

else if (condition₂ holds) **then**

$Q(\text{actions}_3) = \theta_3, \dots$

...

else

$Q(\text{actions}_n) = \theta_n, \dots$


Rule instantiation

- How are the rules applied to the dialogue state?
- The rules are *instantiated* in the Bayesian Network, expanding it with new nodes and dependencies

r₁:

if ($X = \dots \vee Y \neq \dots$) **then**
 $P(V = \dots \wedge W = \dots) = 0.6$

(The ... dots in r_1 should be replaced by concrete values)




Rule instantiation

- The instantiation procedure is similar for utility rules, although one must employ utility and decision nodes:

r₂:

if $(X = \dots \vee Y \neq \dots)$ **then**

$Q(A_1 = \dots \wedge A_2 = \dots) = 3$



Quantification mechanism

- If our domain has a relational structure, the rules must be able to abstract over its *entities*
- To this end, we propose to extend probabilistic rules with a limited form of universal quantification:

$\forall \mathbf{x} = x_1, \dots, x_k:$

if (condition₁(\mathbf{x}) holds) **then**

$P(\text{effect}_1(\mathbf{x})) = \theta_1, P(\text{effect}_2(\mathbf{x})) = \theta_2, \dots$

else if (condition₂(\mathbf{x}) holds) **then**

$P(\text{effect}_3(\mathbf{x})) = \theta_3, \dots$

...

else

$P(\text{effect}_n(\mathbf{x})) = \theta_n, \dots$

Quantification mechanism

- The quantification allows certain variables x_1, \dots, x_k to be *underspecified*.
- The rule will be instantiated for every possible assignment of the underspecified variables.

Example: $\forall o, c :$

if $(a_m = \text{WhatIsColour}(o) \wedge o.\text{colour} = c)$ **then**

$\{P(a'_u = \text{Assert}(\text{Is}(o, c))) = 0.9\}$



The rule will be instantiated for every possible object o and colour c matching the condition



Quantification mechanism

- Why is this quantification mechanism useful?
 - Because it allows the system designer to exploit *high-level abstractions* to encode his domain knowledge
 - Because it is a powerful form of *parameter sharing*, which reduces the number of parameters to estimate... and thereby enables learning algorithms to generalise better and with fewer data



Quantification mechanism

- ... but several questions remain to be addressed (work in progress!)
- Main question: how to keep the formalism tractable?
 - If some variables are underspecified, the algorithm must instantiate the rules for every assignment
 - Need to devise *agressive pruning* techniques to quickly discard irrelevant instantiations

Future evaluation

- The presented framework is being implemented in a dialogue system toolkit called **openDial**
- Evaluation in a human-robot interaction scenario





Conclusion & future work

- We have presented here a simple *quantification mechanism* to augment the expressivity of probabilistic rules
- Such mechanism would enable the rules to directly operate on dialogue states represented as *relational structures*
- Ongoing work on implementation and evaluation