

openDial: A toolkit for building dialogue systems based on probabilistic rules



Pierre Lison University of Oslo, Norway.

Introduction

- Statistical techniques are popular in spoken dialogue systems, but they also present a number of challenges, especially for complex, open-ended domains
- One limitation is that the number of parameters to estimate often grows exponentially with the problem size, and can thus require large amounts of training data
- For most dialogue domains, data is however scarce and expensive to acquire
- One way to address this issue is to rely on more expressive representations, able to capture relevant aspects of the *problem structure* in a compact manner
- We describe here such an abstraction mechanism: **probabilistic rules**
- The rules are specifically devised to encode the kind of structure found in probabilistic models of dialogue (from understanding to management and to generation)

Statistical approaches for spoken dialogue systems

Pros	Cons
Explicit account of uncertainties, increased robustness to errors (e.g. from ASR)	Paucity of appropriate data sets: good domain data is scarce and expensive to acquire!
More natural conversational behaviours, better domain- and user-adaptivity	Scalability to complex domains is a challenge (combinatorial explosion of the state space)

- To test these ideas, we are currently developing a software toolkit called **openDial**
- openDial** employs probabilistic rules as a unifying framework for encoding dialogue processing models and estimating their parameters from interaction data

- Key idea:** exploit *structural knowledge* to yield small, compact probabilistic models.
- Major benefit: the rules are described in exponentially fewer parameters, which are thus easier to learn and generalise to unseen data.
- At runtime, the rules serve as *templates* to instantiate a classical probabilistic model, in the form of a Bayesian Network.
- ... which is then directly used for inference

- Rules take the form of structured mappings from *conditions* to (probabilistic) *effects*:

```

if (condition1 holds) then
  P(effect1) = θ1, P(effect2) = θ2
else if (condition2 holds) then
  P(effect3) = θ3
  
```

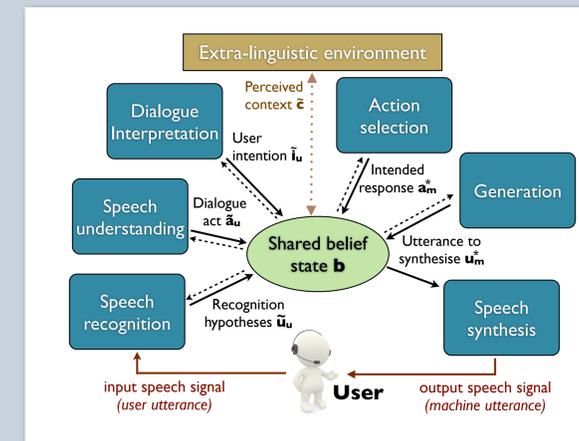
- Conditions are described as logical formulae grounded in a subset of state variables
- Effects are defined similarly, and assign specific values to (new or existing) variables
- For action-selection rules, the effect associates *utilities* to particular actions:

```

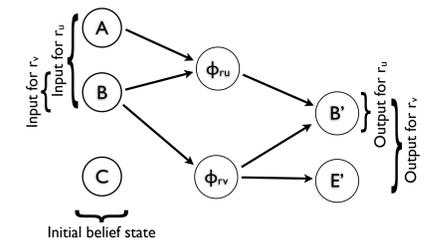
if (condition1 holds) then
  Q(action=value) = θ1
  
```

- Effect probabilities & action utilities are *parameters* which can be estimated from data

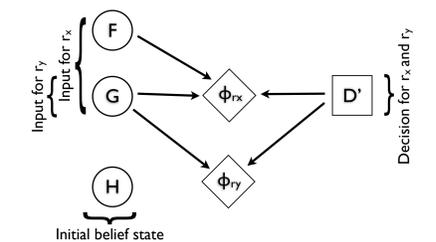
Approach



- The architecture revolves around a shared dialogue state, encoded as a Bayesian Network (belief state b)
- At runtime, this dialogue state is continuously updated via the application of probabilistic rules
- Each processing step (understanding, action selection, generation, etc.) is described by a distinct set of rules
- Decision nodes trigger the execution of specific actions
- External modules such as ASR or TTS can also connect to the shared dialogue state and read/write to it



Example of Bayesian Network updated via probabilistic rules. The initial state contains 3 variables A, B and C, and is expanded with the two rules r_u and r_v . These rules output two variables, B' (updated version of B) and E'.



Another example of Bayesian Network with probabilistic rules, this time used for decision-making. The network shows two rules r_x and r_y specifying the utility of action D'. The dialogue system will then select the value for D' with maximum total utility. The diamond-shaped nodes represent (additive) utilities, while squares denote decision nodes.

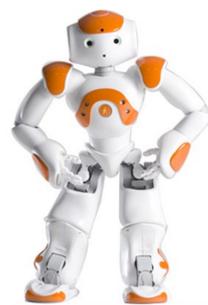
Experiments

- We have evaluated our approach in a human-robot interaction scenario
- The objective was to learn the parameters of the action-selection model (i.e. the dialogue policy) from a small Wizard-of-Oz dataset with 1020 system turns
- Each training sample in the dataset is a pair (dialogue state, action), representing a given (belief) state along with the action selected by the Wizard at that state
- Parameter learning was performed with a Bayesian approach, using a initial prior which gradually narrows down to the values providing the best fit for the data:

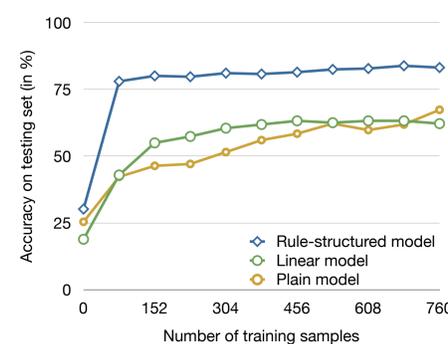
Assuming a training set D , where $d \in D$ is defined as a pair (b_d, t_d) & a collection of rules with parameters θ :

$$\forall d \in D, \forall \theta_i \in \theta, P(\theta_i | d) = \alpha P(t_d | b_d, \theta_i) P(\theta_i)$$

- Baseline: two «flattened» or rolled-out versions of the action selection model, with identical input and output variables, but without the intermediate rule structure
- The empirical results demonstrate that the rule-structured model converges faster and with better generalisation performance



The **Nao** robot was used as experimental platform. The user was instructed to teach the robot a sequence of basic movements (lift the left arm, step forward, kneel down, etc.) using spoken commands.



Learning curve comparing the accuracy of our rule-structured model of action selection against the plain & linear models serving as baselines, plotted as a function of the number of processed training samples from the Wizard-of-Oz data set. We can observe that the rule-structured model is able to converge to near-optimal values after observing only a small fraction of the training set.

Conclusions

- Probabilistic rules used to capture the underlying structure of dialogue models
- Allow developers to exploit powerful generalisations and domain knowledge in the dialogue system design without sacrificing the probabilistic nature of the model
- Very general framework that can express a wide spectrum of models
- In the near future, we aim to extend our approach towards *model-based Bayesian reinforcement learning*, where the parameter distributions are directly updated in a fully online fashion, based on (real or simulated) interaction experience

- On the practical side, we are also developing the **openDial** toolkit, which will enable developers to easily prototype dialogue systems based on probabilistic rules
- The toolkit includes algorithms for efficient inference and parameter estimation, as well as development tools for designing dialogue domains and monitoring interactions