Available online at www.sciencedirect.com

ScienceDirect

Computer Speech and Language 34 (2015) 232–255

COMPUTER SPEECH AND LANGUAGE

www.elsevier.com/locate/csl

# A hybrid approach to dialogue management based on probabilistic rules☆

Pierre Lison *

*Language Technology Group, Department of Informatics, University of Oslo, P.O. Box 1080 Blindern, N-0316 Oslo, Norway*

Received 23 May 2014; received in revised form 1 October 2014; accepted 12 January 2015
Available online 21 January 2015

**Abstract**

We present a new modelling framework for dialogue management based on the concept of *probabilistic rules*. Probabilistic rules are defined as structured mappings between logical conditions and probabilistic effects. They function as high-level templates for probabilistic graphical models and may include unknown parameters whose values are estimated from data using Bayesian inference. Thanks to their use of logical abstractions, probabilistic rules are able to encode the probability and utility models employed in dialogue management in a compact and human-readable form. As a consequence, they can reduce the amount of dialogue data required for parameter estimation and allow system designers to directly incorporate their expert domain knowledge into the dialogue models.

Empirical results of a user evaluation in a human–robot interaction task with 37 participants show that a dialogue manager structured with probabilistic rules outperforms both purely hand-crafted and purely statistical methods on a range of subjective and objective quality metrics. The framework is implemented in a software toolkit called *OpenDial*, which can be used to develop various types of dialogue systems based on probabilistic rules.


*Keywords:* Spoken dialogue systems; Dialogue management; Probabilistic graphical models; Bayesian inference; Human–robot interaction

## 1. Introduction

The design of dialogue strategies is a challenging task in the development of spoken dialogue systems (SDS). The selection of system actions is often grounded in a complex dialogue state encompassing a variety of factors such as the dialogue history, the user goals and preferences, the external context and the task to perform. In addition, spoken dialogue is also riddled with uncertainties arising from speech recognition errors, ambiguous inputs, partially observable environments, and unpredictable dialogue dynamics. These difficulties are particularly striking in the case of human–robot interaction (HRI). By their very definition, human–robot interactions take place in a physical, situated environment that must be captured and monitored by the robotic agent. They must also typically deal with high levels

☆ This paper has been recommended for acceptance by R.K. Moore.

* Tel.: +47 96799812.
*E-mail address:* plison@ifi.uio.no

http://dx.doi.org/10.1016/j.csl.2015.01.001

of noise and uncertainty caused by e.g. imperfect sensors and actuators. The robot's tracking of the current dialogue state is therefore bound to remain partial and error-prone.

Two families of dialogue management approaches have been historically developed to address these issues. The first family relies on hand-crafted strategies, ranging from finite-state automata to more complex inference procedures based on formal logic and classical planning. These strategies provide principled techniques for the interpretation and generation of dialogue moves on the basis of the dialogue participants' mental states (including their shared knowledge). Dialogue is then framed as a collaborative activity in which the interlocutors work together to coordinate their actions, maintain a shared conversational context, resolve open issues and satisfy social obligations (Allen et al., 2000; Larsson, 2002; Jokinen, 2009). Such approaches can yield detailed analyses of various dialogue behaviours, but they generally assume complete observability of the dialogue state and provide only a limited account of errors and uncertainties. In addition, the knowledge bases from which the system's decisions are derived must be completely specified in advance by domain experts. Their deployment in practical applications is thus non-trivial.

The second family relies on statistical modelling techniques (Levin et al., 2000; Roy et al., 2000; Young et al., 2010; Rieser and Lemon, 2011). The dialogue is here represented as a stochastic control process – often a *Markov decision process* (MDP) or a *Partially observable Markov decision process* (POMDP) – and the optimal dialogue strategy is the one that maximises the system's long-term expected utility. These probabilistic models offer an explicit account for the various uncertainties that can arise during the interaction. They also allow the dialogue strategies to be optimised in a data-driven manner instead of relying on hand-crafted mechanisms, making it easier to adapt to new environments or users. However, these probabilistic models typically depend on large amounts of training data to estimate their parameters – a requirement that is hard to satisfy for most dialogue domains. This shortage of relevant datasets is especially critical in human–robot interactions, given the high costs of collecting and annotating dialogue data for these dialogue domains.

This article presents a hybrid approach to dialogue management that seeks to combine the benefits of hand-crafted and statistical techniques in a single framework. As in previous work on POMDPs models for dialogue management, the approach represents the dialogue state as a Bayesian network that is regularly updated with new observations and employed to derive the system's actions. However, the domain models are no longer expressed with traditional factored representations but are instead structured via *probabilistic rules*. As explained in the next pages, the rules can be viewed as high-level templates for probabilistic graphical models. The use of probabilistic rules provides an efficient *abstraction layer* that allows the system designer to capture the domain models in a concise and human-readable form.

The present article is structured as follows. Section 2 reviews the key principles of dialogue management, focusing in particular on MDP- and POMDP-based approaches. Section 3 outlines the formalism of probabilistic rules and their instantiation as nodes of a graphical model. Section 4 describes how the parameters of probabilistic rules can be estimated via Bayesian inference, using either Wizard-of-Oz data (supervised learning) or user interactions (reinforcement learning). Section 5 presents the *OpenDial* toolkit, a domain-independent dialogue toolkit that allows dialogue developers to construct dialogue systems using probabilistic rules. Section 6 describes a user evaluation of this modelling approach in a human–robot interaction domain. Section 7 contrasts the framework with related work. Finally, Section 8 concludes the article and reviews future research directions.

## 2. Dialogue management

### 2.1. System architecture

The general architecture of a spoken dialogue system is depicted in Fig. 1. The user speech signals are first processed by the speech recogniser, resulting in a list of recognition hypotheses $\tilde{u}_u$, where each hypothesis is associated to a particular probability or confidence score.[1] Dialogue understanding then maps these hypotheses into high-level semantic representations of the dialogue act expressed by the user. These dialogue acts are also expressed as a list $\tilde{a}_u$ of semantic hypotheses together with their respective probabilities. Dialogue management is then in charge of selecting the best system action to perform given the current conversational context. The dialogue manager outputs a particular action $a_m$

[1] Throughout this article, we follow the convention of denoting user-specific variables with the subscript $u$ and machine-specific variables with the subscript $m$.