

POMDP-based Statistical Spoken Dialogue Systems: a Review

Steve Young, *Fellow, IEEE*, Milica Gašić, *Member, IEEE*, Blaise Thomson, *Member, IEEE*,
and Jason D Williams, *Member, IEEE*

(Invited Paper)

Abstract—Statistical dialogue systems are motivated by the need for a data-driven framework that reduces the cost of laboriously hand-crafting complex dialogue managers and that provides robustness against the errors created by speech recognisers operating in noisy environments. By including an explicit Bayesian model of uncertainty and by optimising the policy via a reward-driven process, partially observable Markov decision processes (POMDPs) provide such a framework. However, exact model representation and optimisation is computationally intractable. Hence, the practical application of POMDP-based systems requires efficient algorithms and carefully constructed approximations. This review article provides an overview of the current state of the art in the development of POMDP-based spoken dialogue systems.

Index Terms—Spoken dialogue systems, POMDP, reinforcement learning, belief monitoring, policy optimisation.

I. INTRODUCTION

SPOKEN dialogue systems (SDS) allow users to interact with a wide variety of information systems using speech as the primary, and often the only, communication medium [1], [2], [3]. Traditionally, SDS have been mostly deployed in call centre applications where the system can reduce the need for a human operator and thereby reduce costs. More recently, the use of speech interfaces in mobile phones has become common with developments such as Apple’s “Siri” and Nuance’s “Dragon Go!” demonstrating the value of integrating natural, conversational speech interactions into mobile products, applications, and services.

The principal elements of a conventional SDS are shown in Fig 1¹. At each turn t , a spoken language understanding (SLU) component converts each spoken input into an abstract semantic representation called a *user dialogue act* u_t . The system updates its internal state s_t and determines the next system act via a decision rule $a_t = \pi(s_t)$, also known as a *policy*. The system act a_t is then converted back into speech via a natural language generation (NLG) component. The state s_t consists of the variables needed to track the progress of

S. Young, M. Gašić and B. Thomson are with the Dept of Engineering, Trumpington Street, Cambridge University, Cambridge, UK, e-mail: {s.j.y, mg436, brmt2}@eng.cam.ac.uk

JD Williams is with Microsoft Research, Redmond, WA, USA, e-mail: jawillia@microsoft.com

¹Multimodal dialogue is beyond the scope of this paper, but it should be noted that the POMDP framework can be extended to handle multimodal input and output [4]. Depending on the application, both the input and output may include a variety of modalities including gestures, visual displays, haptic feedback, etc. Of course, this could result in larger state spaces, and synchronisation issues would need to be addressed.

the dialogue and the attribute values (often called slots) that determine the user’s requirements. In conventional systems, the policy is usually defined by a flow chart with nodes representing states and actions, and arcs representing user inputs[5], [6].

Despite steady progress over the last few decades in speech recognition technology, the process of converting conversational speech into words still incurs word error rates in the range 15% to 30% in many real world operating environments such as in public spaces and in motor cars [7], [8]. Systems which interpret and respond to spoken commands must therefore implement dialogue strategies that account for the unreliability of the input and provide error checking and recovery mechanisms. As a consequence, conventional deterministic flowchart-based systems are expensive to build and often fragile in operation.




Fig. 1. Components of a finite state-based spoken dialogue system. At each turn the input speech is converted to an abstract representation of the user’s intent u_t , the dialogue state s_t is updated and a deterministic decision rule called a *policy* maps the state into an action a_t in response.

During the last few years, a new approach to dialogue management has emerged based on the mathematical framework of partially observable Markov decision processes (POMDPs²) [9], [10], [11]. This approach assumes that dialogue evolves as a Markov process, i.e., starting in some initial state s_0 , each subsequent state is modelled by a transition probability: $p(s_t|s_{t-1}, a_{t-1})$. The state s_t is not directly observable reflecting the uncertainty in the interpretation of user utterances; instead, at each turn, the system regards the output of the SLU as a noisy observation o_t of the user input with probability $p(o_t|s_t)$ (see Fig 2). The transition and observation probability functions are represented by a suitable stochastic model, called here the *dialogue model* \mathcal{M} . The decision as to which action to take at each turn is determined by a second stochastic

²pronounced “pom dee pees”

model encoding the policy \mathcal{P} . As the dialogue progresses, a reward is assigned at each step designed to mirror the desired characteristics of the dialogue system. The dialogue model \mathcal{M} and policy model \mathcal{P} can then be optimised by maximising the expected accumulated sum of these rewards either on-line through interaction with users or off-line from a corpus of dialogues collected within a similar domain.




Fig. 2. Components of a POMDP-based spoken dialogue system. In contrast to Fig. 1, the decoded input speech is now regarded as a noisy observation o_t of the underlying user intent u_t . Since u_t is hidden, the system maintains a distribution b_t over all possible dialogue states and instead of trying to estimate the hidden dialogue state, the system response is determined directly from b_t . In addition, the dialogue model and policy are parameterised, and given an appropriate reward function, they can be optimised using reinforcement learning.

This POMDP-based model of dialogue combines two key ideas: belief state tracking and reinforcement learning. These ideas are separable and have benefits on their own. However, combining them results in a complete and well-founded mathematical framework that offers opportunities for further synergistic gains. The potential advantages of this approach compared to conventional methods can be summarised as follows:

- 1) the belief state provides an explicit representation of uncertainty leading to systems that are much more robust to speech recognition errors [11]. The posterior probability of the belief state after each user input is updated via Bayesian inference in a process called *belief monitoring*. The design of the belief state allows user behaviour to be captured via the model priors and the inference process is able to exploit the full distribution of recognition hypotheses such as confusion networks and N -best lists. Thus, evidence is integrated across turns such that a single error has significantly reduced impact, and in contrast to conventional systems, user persistence is rewarded. If the user repeats something often enough, the system's belief in what they said will increase in time as long as the correct hypothesis appears repeatedly in the N -best list.
- 2) by maintaining a belief distribution over all states, the system is effectively pursuing all possible dialogue paths in parallel, choosing its next action not based on the most likely state but on the probability distribution across all states. When the user signals a problem, the probability of the current most likely state is reduced and the focus simply switches to another state. Thus, there is no requirement for back-tracking or specific error correction dialogues. This allows powerful dialogue policies to be

embedded in a simple homogenous mapping from belief state to action.

- 3) the explicit representation of state and policy-derived action allows dialogue design criteria to be incorporated by associating rewards with state-action pairs. The sum of these rewards constitutes an objective measure of dialogue performance and enables reinforcement learning to be used to maximise performance both off-line using dialogue corpora and on-line through interaction with real users. This leads to optimal decision policies, avoids the cost of expensive manual tuning and refinement procedures, and enables more complex planning to be implemented than would be feasible using only manual hand-crafted designs.

Converting these potential benefits of the POMDP approach into practice is, however, far from trivial and there are many issues to resolve. The state-action space of a real-world SDS is extremely large and its efficient representation and manipulation requires complex algorithms and software. Real-time Bayesian inference is equally challenging and exact policy learning for POMDPs is intractable, hence efficient approximation techniques must be used. Finally, the most straightforward way to optimise a POMDP-based SDS is through direct interaction with users. However, real users willing to help train a system are often not available in sufficient numbers, hence, user simulators are required that can replicate user behaviour with sufficient accuracy to enable model parameters to be optimised to an acceptable level of performance (a discussion of user simulation is given in section V).

Despite these difficulties considerable progress has been made over the last five years in solving these problems and the purpose of this paper is to review that progress and provide a coherent up-to-date view of the state of the art in POMDP-based dialogue systems.

The paper is organised as follows. Firstly, section II outlines the basic mathematics underlying POMDP-based dialogue systems in order to provide the necessary background for the subsequent sections and to establish a consistent notation. Section III then explains the options available for approximating the belief state and presents algorithms for efficient belief monitoring. Section IV reviews policy representation, the use of reward functions and policy optimisation via reinforcement learning. Section V completes the review of core technology with an overview of current approaches to user simulation. Having established the basic framework, sections VI and VII review a number of recent developments in optimisation of dialogue model parameters and fast adaptation. To give some indication of the potential for real-world deployment, section VIII briefly describes some existing prototype systems and applications that incorporate POMDP-based dialogue management, with some example evaluations given in section IX. Finally, for completeness, section X provides a historical perspective and section XI concludes.

II. PARTIALLY OBSERVABLE MARKOV DECISION PROCESSES

Formally, a partially observable Markov decision process is defined as a tuple $(\mathbb{S}, \mathbb{A}, \mathbb{T}, \mathbb{R}, \mathbb{O}, \mathbb{Z}, \gamma, b_0)$ where \mathbb{S} is a set of states with $s \in \mathbb{S}$; \mathbb{A} is a set of actions with $a \in \mathbb{A}$; \mathbb{T} defines a transition probability $P(s_t|s_{t-1}, a_{t-1})$; \mathbb{R} defines the expected (immediate, real-valued) reward $r(s_t, a_t) \in \mathbb{R}$; \mathbb{O} is a set of observations with $o \in \mathbb{O}$; \mathbb{Z} defines an observation probability $P(o_t|s_t, a_{t-1})$; γ is a geometric discount factor $0 \leq \gamma \leq 1$; and b_0 is an initial belief state, defined below.

The POMDP operates as follows. At each time-step, the world is in some unobserved state s_t . Since s_t is not known exactly, a *distribution* over possible states called a *belief state* b_t is maintained where $b_t(s_t)$ indicates the probability of being in a particular state s_t . Based on b_t , the machine selects an action a_t , receives a reward r_t , and transitions to (unobserved) state s_{t+1} , where s_{t+1} depends only on s_t and a_t . The machine then receives an observation o_{t+1} , which is dependent on s_{t+1} and a_t . This process is represented graphically as an influence diagram in Fig 3.




Fig. 3. A partially observable Markov decision process, shown as an influence diagram. In this paper, clear circles are hidden random variables; shaded circles are observed random variables; squares are system actions; diamonds are real-valued rewards; and arrows show causality.

Given an existing belief state b_t , the last system action a_t , and a new observation o_{t+1} , the new updated belief state b_{t+1} is given by [12]:

$$b_{t+1}(s_{t+1}) = \eta P(o_{t+1}|s_{t+1}, a_t) \sum_{s_t} P(s_{t+1}|s_t, a_t) b_t(s_t) \quad (1)$$

where $\eta = P(o_{t+1}|b_t, a_t)$ is a normalization constant and where b_0 is the initial belief state distribution before the first system action has been taken.³

The system action is determined by a policy π , which can be represented in a variety of ways. It is most commonly either a deterministic mapping from belief states to actions $\pi(b) \in \mathbb{A}$ or stochastically via a distribution over actions $\pi(a|b) \in [0, 1]$ where $\pi(a|b)$ is the probability of taking action a in belief state b , and $\sum_a \pi(a|b) = 1 \forall b$. For convenience,

³The notation for belief states can be confusing: b_t represents a probability distribution over the hidden state space \mathbb{S} at time t ; $b_t(s)$ denotes the probability of a specific state s given belief state b_t ; $b_{t+1}(s)$ represents the probability of state s given a new updated belief state b_{t+1} , which in general will be different from $b_t(s)$.

both types of policy will use the same symbol π , with the presence of the action in the notation determining whether the policy is deterministic or stochastic. Note, however, that other definitions are possible such as finite state controllers [13], or mappings from finite-length sequences of observations to actions (c.f. predictive state representations [14]).

The discounted sum of rewards expected by starting in belief state b_t and following policy π is given by the *value function* $V^\pi(b_t) = \mathbb{E}[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots]$, which can be expressed recursively for a deterministic policy as

$$V^\pi(b_t) = r(b_t, \pi(b_t)) + \gamma \sum_{o_{t+1}} P(o_{t+1}|b_t, \pi(b_t)) V^\pi(b_{t+1}) \quad (2)$$

and by

$$V^\pi(b_t) = \sum_{a_t} \pi(a_t|b_t) \left\{ r(b_t, a_t) + \gamma \sum_{o_{t+1}} P(o_{t+1}|b_t, a_t) V^\pi(b_{t+1}) \right\} \quad (3)$$

for a stochastic policy. A related quantity is the *Q-function* $Q^\pi(b, a)$, which provides the expected discounted sum of rewards if a specific action a is taken given belief state b , and then policy π is followed. Clearly for a deterministic policy, $V^\pi(b) = Q^\pi(b, \pi(b))$ and for a stochastic policy

$$V^\pi(b) = \sum_a \pi(a|b) Q^\pi(b, a). \quad (4)$$

An *optimal policy* π^* is one that maximizes V^π to yield V^*

$$V^*(b_t) = \max_{a_t} \left[r(b_t, a_t) + \gamma \sum_{o_{t+1}} P(o_{t+1}|b_t, a_t) V^*(b_{t+1}) \right] \quad (5)$$

which is the Bellman optimality equation for POMDPs [15]. In the POMDP literature, finding a policy π that satisfies (5) is often called “solving” or “optimizing” the POMDP. For simple tasks, both exact [12] and approximate [16], [17], [18], [19], [20] solution methods have been developed. However, standard POMDP methods do not scale to the complexity needed to represent a real-world dialogue system. Even in a moderately-sized system, the number of states, actions, and observations can each easily be more than 10^{10} . Even enumerating $P(s_{t+1}|s_t, a_t)$ is intractable, and as a result computing (1) directly and applying direct solution methods to (5) is very difficult. Instead, approximations have been developed that exploit domain-specific properties of the spoken dialogue task in order to provide compact representations for both the model and the policy; and to allow tractable algorithms for performing belief monitoring and policy optimisation. These are described in the following sections.

III. BELIEF STATE REPRESENTATION AND MONITORING

This section reviews the possible approaches to representing the dialogue model \mathcal{M} shown in Fig. 2. In a practical task-oriented SDS, the state must encode three distinct types of

information: the user's goal g_t , the intent of the most recent user utterance u_t and the dialogue history h_t [21], [11]. The goal encompasses the information that must be gleaned from the user in order to fulfil the task, the most recent user utterance represents what was actually said in contrast to what was recognised, and the history tracks pertinent information relating to previous turns. This suggests that the state should be factored into three components

$$s_t = (g_t, u_t, h_t). \quad (6)$$

The resulting influence diagram is shown in Fig. 4 in which some reasonable independence assumptions have been introduced. Factoring the state in this way is helpful because it reduces the dimensions of the state transition matrix and it reduces the number of conditional dependencies.




Fig. 4. Influence diagram representation of a factored state SDS-POMDP. The hidden dialogue state s_t is factored into a user goal g_t , the user's last input u_t and key elements of the dialogue history h_t . This allows conditional independence assumptions to be introduced both within the state and from one time-slice to the next. No arc is required from actions to observations since the recognised speech o_t is conditionally independent given the user's actual utterance u_t .

Plugging the factorisation in (6) into the belief update equation (1) and simplifying according to the independence assumptions shown in Fig 4 gives the basic belief update equation for a statistical SDS:

$$b_{t+1}(g_{t+1}, u_{t+1}, h_{t+1}) = \eta P(o_{t+1}|u_{t+1}) \quad (a)$$

$$\cdot P(u_{t+1}|g_{t+1}, a_t) \quad (b)$$

$$\cdot \sum_{g_t} P(g_{t+1}|g_t, a_t) \quad (c)$$

$$\cdot \sum_{h_t} P(h_{t+1}|g_{t+1}, u_{t+1}, h_t, a_t) \quad (d)$$

$$\cdot b_t(g_t, h_t) \quad (7)$$

The terms on the right hand side in (7) reflect each of the factors determining the belief state and the underlying models which are therefore needed to represent these factors in a practical system:

(a) the *observation model* represents the probability of an observation o given the user's actual utterance u . This encapsulates the effects of speech understanding errors.

- (b) the *user model* represents the likelihood that the user would utter u given the previous system output and the new system state. This encapsulates user behaviour.
- (c) the *goal transition model* represents the likelihood that the user goal has changed.
- (d) the *history model* represents the system's memory of the dialogue to date.

There are of course possible variations to this factorisation. For example, user affect may also be factorised out [22] but most current approaches fit broadly within this model.

While the factorisation in (6) does significantly reduce the POMDP model complexity, it is still too complex to support tractable real-world systems. Further approximation is therefore necessary for which two main approaches have emerged:

- 1) the *N-best* approach including pruning and recombination strategies [23], [24], [25], [26], and
- 2) the factored Bayesian Network approach [27], [22], [28].

These two approaches are discussed below.

A. *N-best* approaches

In *N-best* approaches, the belief state is approximated by a list of the most likely states (or groups of states) with their probabilities. This means that dialogue states corresponding to the most likely interpretations of the user's intent are well modelled, with other states given low probability mass. One example of this approach is the Hidden Information State model (HIS) [23], which groups similar user goals into equivalence classes called *partitions* on the assumption that all of the goals in the same partition are equally probable. The partitions are tree-structured to take account of the dependencies defined in the domain ontology and they are built using slot-value pairs from the *N-best* list of recognition hypotheses and the last system output. The combination of a partition, a user act from the *N-best* list and the associated dialogue history forms a *hypothesis*. A probability distribution over the most likely hypotheses is maintained during the dialogue and this constitutes the belief space. Belief monitoring then requires only the hypothesis beliefs to be updated and since there are relatively few hypotheses, this can easily be done in real time. The update equation for the HIS model follows directly from (7) with the further simplification that the user goal is normally assumed to be constant [23]

$$b_{t+1}(p_{t+1}, u_{t+1}, h_{t+1}) = \eta P(o_{t+1}|u_{t+1}) P(u_{t+1}|p_{t+1}, a_t) \quad (8)$$

$$\sum_{h_t} P(h_{t+1}|p_{t+1}, u_{t+1}, h_t, a_t) P(p_{t+1}|p_t) b(h_t)$$

where p_t is a partition and the term $P(p_{t+1}|p_t)$ represents the probability that a partition p_t will be split into two sub-partitions $p_t \rightarrow \{p_{t+1}, p_t - p_{t+1}\}$. A similar approach is taken in [24], except that both slot values and their complements are used to build a frame. This is particularly useful in negotiation-type dialogues, where the users can change their mind and ask for an alternative. It also enables a form of first-order logic to be expressed and understood by the system. However, the

domain ontology is not used when building frames losing the benefits of modelling conditional dependence. In [26], it is shown that complements can also be incorporated into the HIS approach giving the advantage of modelling both conditional dependence and first order logic.

The N -best approach can also be viewed as running N conventional dialogue systems in parallel such that each parallel thread tracks a different interpretation of what the user said. In this case, a list of hypothesised dialogue states is maintained; associated probabilities can be assigned using a mixture distribution [29], a discriminative classifier [30], or with a heuristic scoring function [31]. Although it lacks the compactness of the HIS representation, this approach provides a simple upgrade path for conventional deterministic dialogue systems.

Both the partition-based and the frame-based approaches have the inherent problem that the number of partitions (frames) grows exponentially with dialogue length. Hence, in practical implementations, some form of pruning is required. In [25], partitions are organised into a tree and then low probability partitions are recombined with their parent, thus limiting the number of partitions. This has the problem however that high probability slots can be pruned if they are distributed over many partitions. A more effective solution is proposed in [26], where at every dialogue turn a marginal probability distribution is calculated for each slot. Then, low probability slot-value pairs are pruned by recombining all the partitions that have that slot-value pair with the partitions that have the complement of that slot-value pair. In this way, the method supports dialogues of arbitrary length.

B. Factored approaches

The alternative to maintaining an N -best list of user goals is to factor the user goal into concepts that can be spoken about by the system. This is illustrated in Fig. 5, which shows an example from a tourist information system in which entities have a *type* such as $\{hotel, restaurant, bar\}$, the kind of *food* served and an *area* such as $\{north, south, centre, \dots\}$ [28]. The *food* and *area* slots are dependent only on the type, even though in practice, there are many more restaurants in the centre of town. This illustrates the differing trade-offs between the N -best approach, which can model all dependencies but with an incomplete distribution, and the slot-level factoring approach, which can handle only a limited number of dependencies but can model the complete distribution.

Once the Bayesian network for the slot-level factoring has been compiled, standard belief propagation algorithms can be used to update the beliefs [32]. In the case of concepts that are conditionally independent, the belief propagation algorithm will give exact updates of the marginal distributions for each of the concepts, as used in [22]. Limited dependencies can also be modelled, although this requires approximate methods such as loopy belief propagation. The Bayesian update of dialogue state (BUDS) model is one example that is based on this approach [27], particle filters can also be used [33].

It is interesting to note that N -best approximations can also be used within a factored model. Instead of the factored model

updating beliefs for all the possible values of a node in the network, only an N -best list of values is updated. Partitions of the states in a node are also possible [34]. This combines many of the benefits of the two approaches.

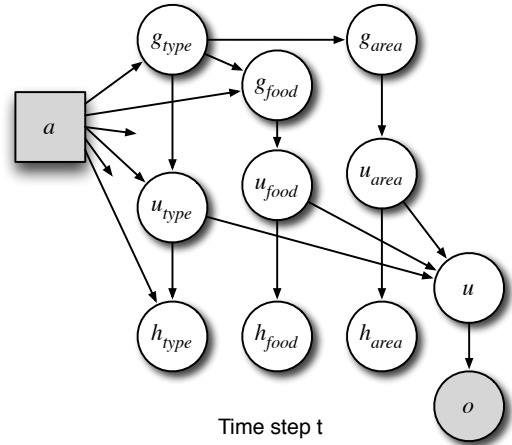


Fig. 5. Influence diagram for a single time slice of a BUDS POMDP in which the state is further factored into concept or slot level components. In this simplified example taken from the tourist information domain, three slot level values are required: the type of venue, the kind of food served and the area in which the venue is located. Note that food and area are conditionally independent given the type of venue and last action. (Although not shown explicitly to avoid cluttering the diagram, all slot-level nodes are dependent on the last system action.)

IV. POLICY REPRESENTATION AND REINFORCEMENT LEARNING

This section presents the representation and estimation of the policy model \mathcal{P} shown in Fig. 2, which provides a mapping between the belief state b and the appropriate system action a . The objective is to find an optimal policy π^* that maximizes the expected sum of discounted rewards at the end of the dialogue.

The belief space of a POMDP spans an $(|\mathcal{S}|-1)$ -dimensional simplex where $|\mathcal{S}|$ is the cardinality of the underlying hidden state space. Points in belief space that are close to each other will share the same action, and hence, a non-parametric policy must encode firstly a partitioning of belief space such that all points within any partition map to the same action, and secondly it must encode the optimal action to take for each partition. Whilst exact representations of a POMDP dialogue policy are possible, for example by compressing belief space [35] or dynamically re-assigning states [36], exact representations are all intractable for real world problems such as spoken dialogue systems where both the state and action spaces will typically be very large. Thus, a compact representation of the policy is essential.

Fortunately, there are some mitigating constraints that can be exploited. Firstly, only a relatively small part of belief space will actually be visited during any normal dialogue, and secondly, the range of plausible actions at any specific point in belief space will often be restricted. This introduces the notion of a compressed feature space called *summary space* in which both states and actions are simplified in order to

allow tractable policy representation and optimisation [37], [38]. Summary space can therefore be regarded as a subspace of the full master space whereby belief tracking is performed in master space, and decision-taking and policy optimisation take place in summary space. The run-time operation of a master-summary space POMDP is therefore as follows. After belief updating, the belief state b in master space is mapped to a vector of features \hat{b} and a corresponding set of candidate actions $\{\hat{a}\}$. The policy is then used to select the best action to take $\hat{b} \rightarrow \hat{a}$ from the set of candidate actions and a second heuristic is used to map \hat{a} back into a full action a in master space.

Summary-space mapping requires two components: a mechanism to select candidate actions in master space and functions to extract features from the belief state and candidate actions. The simplest method for selecting candidate actions is to include all types of *dialog acts* (e.g., *greet*, *ask*, *confirm*, *inform*, etc.) applicable to any *concept* or *slot* (e.g., *venue-type*, *food-type*, *star-rating*, etc.), populating slot values by highest belief [28], [37]. This approach has the benefit of being entirely automatic; however, it still admits some spurious candidate actions, such as taking a greeting action in the middle of a dialog, or attempting to confirm a value before the system has asked about it. An alternative method for selecting candidate actions is to construct a hand-crafted *partial program* [39], [40] or a Markov logic network [41]. These approaches have the benefit of allowing arbitrary human knowledge about dialogue flow to be incorporated, and to explicitly set business rules, for example, requiring that certain actions be taken before others such as successfully collecting a password before allowing funds to be transferred. It has also been shown that constraining the set of candidate actions results in faster convergence to an optimal policy, since many spurious actions are pruned away [40]. However, rules require human effort to encode, and there is a risk that optimal candidate actions may be inadvertently excluded. Intermediate methods are possible too, such as allowing every dialogue act to be a candidate action, but constraining the slots that the act operates on using hand-crafted heuristics [42].

The second component required for summary-space mapping are functions to extract features from the belief state and each candidate action. For actions, typically one binary feature is created for each dialog act, or for each valid dialog act/slot pair, such as *confirm(food)*. This generally results in a 20-30 dimensional vector of action features, where each dimension represents a unique (summary) action. State features are typically heterogeneous, consisting of real-valued quantities, binary values, categorical values, etc. Typical state features include: the belief in the top N user goals or partitions; the top marginal belief in each slot; properties of the top user goal or partition, such as the number of matching database entries; an indication of which system actions are available; properties of the dialogue history, such as whether the top user goal/partition has been confirmed; the most likely previous user action; or combinations of these features [37], [40], [28], [42], [43]. A typical system has 5 to 25 features in total which are usually hand-selected, although some work has been done to automate feature selection [43], [44]. State features

need not be limited to information in the belief state: features may also draw on information being tracked outside the belief state, such as information in databases, information from past dialogues, usage context, etc.

Given a specific summary space, a policy may be represented as an explicit deterministic mapping: $\pi(\hat{b}) \rightarrow \hat{a}$ or as a conditional probability distribution $\pi(\hat{b}, \hat{a}) = p(\hat{a}|\hat{b})$ where the required action is selected by sampling the distribution. Note that the policy is now a function of the summary belief state and actions, instead of the original belief state and actions. One can think of this either as a new function that gives an approximation to the value in the original space, or as a policy over a new Markov decision process, where the states and actions are now the summary states and actions. The same is true of the Q function, and both these functions will now be used in this context.

In the case of deterministic mappings, the dominant approach is to find an optimal Q -function (see section II) Q^* that maximises $\max_{\hat{a}}\{Q(\hat{b}, \hat{a})\}$ for all \hat{b} , then

$$\pi^*(\hat{b}) = \arg \max_{\hat{a}}\{Q^*(\hat{b}, \hat{a})\}. \quad (9)$$

The Q -function itself can either be parametric or non-parametric in which case the belief state is quantised into a codebook $\{\hat{b}_i\}$ and $Q^*(\hat{b}_i, \hat{a})$ can be computed for each discrete codebook entry \hat{b}_i .

To illustrate some of these options, five different methods of policy optimisation are now presented: planning under uncertainty, value iteration, Monte-Carlo optimisation, least-squares policy iteration, and natural actor-critic. These methods have been selected because all have been applied to end-to-end working dialogue systems. Some further optimisation methods will be covered later in section VII on fast adaptation. While representative, these five methods are not exhaustive – for example, Q-Learning [45], [46], [47], [48] and SARSA [49], [50], [51] are also popular techniques.

A. Planning under uncertainty

One approach to policy optimisation is to view action selection in summary space as *planning under uncertainty*. In this approach, \hat{b} is viewed not as a vector of arbitrary features, but rather as a distribution over a hidden state variable in summary space \hat{s} . This hidden summary state might express, for example, whether the top hypothesis in master space is correct or not. In addition, it is assumed observations o (in master space) can be mapped to a compact observation in summary space \hat{o} .

With these definitions in place, dialogues can be collected with a random policy, and three models of summary space dynamics can be estimated – $P(\hat{s}'|\hat{s}, \hat{a})$, $P(\hat{o}'|\hat{s}, \hat{a})$, and $R(\hat{s}, \hat{a})$. These models define a classical POMDP *in summary space* and admit standard methods for planning under uncertainty, including different versions of point-based value iteration [16], [18], SARSOP [52], heuristic search value iteration [17], [19], and short look-aheads [53], [22]. These methods have been applied to dialogue systems in a variety of domains, mainly in simulation [37], [54], [53], [22], [55], [56], [57], [58], [59], [60] and also in an end-to-end spoken dialogue system [61].

Optimization by planning under uncertainty is attractive in that optimization methods are principled, and in some cases provide well-defined optimality bounds [17], [19]. Another strength is that a policy can be computed from any corpus that contains sufficient exploration of possible dialog paths. However, it suffers from several problems. First, the summary state \hat{s} and observation \hat{o} may be difficult to define, and incorporation of extrinsic features in \hat{b} may require careful engineering of the models $P(\hat{s}'|\hat{s}, \hat{a})$ and $P(\hat{o}'|\hat{s}, \hat{a})$. Also, while great strides have been made in scaling optimization methods for planning under uncertainty, supporting large state or observation spaces in summary space can be problematic. For these reasons, techniques have been developed which avoid defining states and observations in summary space. These are described in the next four sections.

B. Value Iteration

Instead of defining states and observations in summary space, in *value iteration* belief states \hat{b} are viewed as arbitrary feature vectors, and dynamics are estimated directly in this feature space. Starting again with dialogues collected with actions taken randomly, feature vectors are quantised into grid points $\{\hat{b}_i\}$, and then the transition function $P(\hat{b}_j|\hat{b}_i, \hat{a})$ and reward function $R(\hat{b}_i, \hat{a})$ are estimated [40], [62]. Finally, standard value iteration can be applied [63]

$$Q(\hat{b}_i, \hat{a}) = r(\hat{b}_i, \hat{a}) + \gamma \sum_j P(\hat{b}_j|\hat{b}_i, \hat{a}) \max_{\hat{a}'} Q(\hat{b}_j, \hat{a}') \quad (10)$$

where Eq. 10 is applied repeatedly until convergence. The clusters \hat{b}_i can either be distinct *hard* clusters, or *soft* clusters where observed \hat{b} are shared by nearby \hat{b}_i . In this case, the contributions of \hat{b}_i that are further away from \hat{b} can be down-weighted [40], [62].

Value iteration's main attraction is that it is simple to apply, and it allows a policy to be estimated from any corpus that contains sufficient exploration of possible dialogue paths. However, it suffers from several problems. Firstly, it requires an estimate of transition dynamics for the entire space $P(\hat{b}_j|\hat{b}_i, \hat{a})$, when in practice it is only important to estimate the portions of space traversed by the currently optimal policy. Hence, value iteration is sample inefficient. Secondly, if the state features are not Markovian – i.e., if they do not accurately capture all relevant history for conditioning transitions – errors will be introduced into the policy. To address both of these issues, incremental *on-line* learning methods have also been explored.

C. Monte Carlo optimization

In Monte Carlo optimization, the value function Q is estimated *on-line* – iteratively, whilst interacting with the user (or more likely interacting with a user simulator - see section V). The current estimate of the policy guides future action selection, so less exploration time is expended on regions of the state space with low value. In addition, updates are performed at the end of each dialogue; hence, whole-dialogue rewards can be used to directly value the current policy, which

mitigates against the effects of any non-Markovian dynamics in the state features.

To perform Monte Carlo optimization, the sequence of states, actions, and rewards observed in each dialogue of length T are recorded as tuples $(\hat{b}_{i(t)}, \hat{a}_t, r_t)$ for $t = 0, 1, \dots, T$. The discounted return from each belief point \hat{b}_i visited at time t is given by

$$R_{i(t)} = \sum_{t \leq \tau \leq T} \gamma^{(\tau-t)} r_\tau. \quad (11)$$

Two functions accumulate sums across all dialogues in the corpus. First, $q(i, \hat{a})$ is a sum of all returns R observed when action \hat{a} was taken in grid point \hat{b}_i . Second, $n(i, \hat{a})$ is a count of the instances of action \hat{a} being taken in gridpoint \hat{b}_i . Both $q(i, \hat{a})$ and $n(i, \hat{a})$ are initialized to zero for all i and \hat{a} . At the end of each dialogue, these tallies are then updated

$$q(i(t), \hat{a}) = q(i(t), \hat{a}) + R_{i(t)} \quad (12)$$

$$n(i(t), \hat{a}) = n(i(t), \hat{a}) + 1, \quad (13)$$

for $t = 0, 1, \dots, T$. The Q values can then be re-estimated for all belief points i visited in the dialogue and all \hat{a} as

$$Q(\hat{b}_i, \hat{a}) = \frac{1}{n(i, \hat{a})} q(i, \hat{a}). \quad (14)$$

To converge to an optimal policy, the dialogue system must take a mixture of currently optimal and exploratory actions

$$\hat{a} = \begin{cases} \text{RandomAction} & \text{with probability } \epsilon \\ \arg \max_{\hat{a}} Q(\hat{b}_i, \hat{a}) & \text{with probability } 1 - \epsilon \end{cases} \quad (15)$$

where ϵ controls the exploration rate. Typically this is large initially but reduces as learning progresses. As in the value iteration case, the state discretization can be done by clustering neighbouring points into distinct *hard* clusters [42], or *soft* clusters where the contribution of a belief point \hat{b} to the tallies $q(i, \hat{a})$ and $n(i, \hat{a})$ is reduced in proportion to its distance from the template point \hat{b}_i [64].

One important consideration for Monte Carlo optimisation is the number of grid points to use. If too many grid points are used, poor performance will result due to over-fitting. Conversely, having too few grid points leads to poor performance due to a lack of discrimination in decision making. For building practical systems with 5 summary state components and 10 distinct summary actions, 750–1000 grid points have been found to be sufficient when trained with around 100,000 simulated dialogues [42].

As compared to value iteration, on-line Monte Carlo optimisation is more sample efficient and suffers less when the transitions are non-Markovian in the state features. However, since Monte Carlo optimisation is an on-line method, it cannot be performed using a previously collected corpus. Like value iteration, Monte Carlo optimization requires quantizing the feature space into grid regions. As the number of features grows, this becomes a source of intractability, which motivates the use of functional approximations for the Q -function, presented next.

D. Least-squares policy iteration

Instead of using discrete grid points in summary space, the next two methods utilise linear models to represent a policy. The first is *least-squares policy iteration* (LSPI) [65]. LSPI assumes that the Q function can be represented as a weighted sum of the features ϕ

$$Q(b, \hat{a}) \approx \phi(b, \hat{a})^\top \theta \quad (16)$$

where θ is a (column) weight vector, $\phi(b, \hat{a})$ outputs features of b that are important for making decisions relevant to summary action \hat{a} , and θ and $\phi(b, \hat{a})$ are both of size K . Notice that the functional mapping from master belief space to summary belief space is now explicit in the function ϕ . The Q -function is thus approximated as a plane in feature space. This may result in sub-optimal policies if the representation is not rich enough.

LSPI operates on a given, *fixed* dialogue corpus with $N + 1$ turns $\{b_n, \hat{a}_n, r_n, b'_n\}$ where b'_n is the belief state following b_n in the corpus. LSPI begins with some policy π , and optimizes it by iteratively estimating its value function Q , then using Q to improve π . To estimate Q , we start with the identity that defines the value function Q

$$Q(b, \hat{a}) = r + \gamma Q(b', \pi(\hat{a})). \quad (17)$$

Substituting Eq 16 into Eq 17 yields a system of N equations in K unknowns

$$\phi(b_n, \hat{a}_n)^\top \theta = r_n + \gamma \phi(b'_n, \pi(b'_n))^\top \theta \quad (18)$$

for $n = 0..N-1$. This system of equations can be solved in a least squares sense by setting $\theta = Z^{-1}y$ where [65]

$$Z = \sum_n \phi(b_n, \hat{a}_n) (\phi(b_n, \hat{a}_n) - \gamma \phi(b'_n, \pi(b'_n)))^\top \quad (19)$$

$$y = \sum_n \phi(b_n, \hat{a}_n) r_n \quad (20)$$

where Z is a $K \times K$ matrix and y is a K -dimensional vector. $\pi(b)$ is then updated by

$$\pi(b) = \arg \max_{\hat{a}} \phi(b, \hat{a})^\top \theta \quad (21)$$

and re-estimation of θ and π then continues until both have converged.

LSPI is attractive in that it avoids estimating grid points in belief space in favour of a linear model for Q . It also has no learning parameters to tune and it can be used both on-line and off-line, although in the case of the latter it can be difficult to find corpora with sufficient variability to learn policies that are significantly different to the policy used to collect the corpus.

One drawback of LSPI is that the least-squares solution step is cubic in K ; hence, for practical applications of LSPI the feature set must be kept quite small. To assist this, methods have been developed that perform feature selection in conjunction with LSPI [43] whereby, at each iteration, a fast but approximate estimate of θ is computed on the full set of features; then only the features with the highest weights are used in the expensive but more accurate least-squares solution. Experiments show this method produces policies of similar performance but with far fewer features [43].

E. Natural actor-critic optimization

LSPI uses a linear model to estimate the value function Q . An alternative approach is to represent the policy directly as an explicit probability distribution over actions, parameterized by an adjustable weight vector θ ,

$$\pi(\hat{a}|b, \theta) = \frac{e^{\theta \cdot \phi_{\hat{a}}(b)}}{\sum_{\hat{a}'} e^{\theta \cdot \phi_{\hat{a}'}(b)}}, \quad (22)$$

where $\phi_{\hat{a}}(b)$ determines the features of the belief state b that are important for making decisions relevant to summary action \hat{a} . As above, ϕ is specified by the system designer, and ϕ and θ both have dimensionality K . π is as before, with the θ simply indexing which policy parameters to use (i.e., policy $\pi_\theta(\hat{a}|b) = \pi(\hat{a}|b, \theta)$).

Representing the policy in this way has two advantages. Firstly, it avoids the need to choose grid points and the problems that result. Secondly, since $\pi(\hat{a}|b, \theta)$ is differentiable with respect to θ , it allows gradient ascent to be applied for optimisation. The expected cumulative reward for a single dialogue as a function of θ is

$$R(\theta) = \mathbf{E} \left[\sum_{t=0}^{T-1} r(b_t, \hat{a}_t) | \pi_\theta \right]. \quad (23)$$

By rewriting the log likelihood-ratio [66] and Monte Carlo sampling from a batch of N dialogues, the gradient can be estimated as

$$\nabla R(\theta) = \frac{1}{N} \sum_{n=1}^N \sum_{t=0}^{T_n-1} \nabla_\theta \log \pi(\hat{a}_t^n | b_t^n, \theta) Q(b_t^n, \hat{a}_t^n) \quad (24)$$

Although (24) provides an estimate for the plain gradient, it has been shown that the natural gradient $\tilde{\nabla} R(\theta) = F_\theta^{-1} \nabla R(\theta)$ is more effective for optimisation of statistical models where F_θ is the Fisher Information Matrix [67]. Furthermore, this natural gradient w can be found without actually computing the Fisher Matrix by using a least square method to solve the following system of equations

$$R^n = \left[\sum_{t=0}^{T_n-1} \nabla_\theta \log \pi(\hat{a}_t^n | b_t^n, \theta)^\top \right] \cdot w + C \quad (25)$$

for each dialogue $n = 1..N$ where R^n is the actual cumulative reward earned for dialogue n and C is a constant that can be interpreted as the expected cumulative reward of the dialogue system starting from the initial state. The total expected reward can then be iteratively maximised through gradient ascent via the parameter update $\theta' \leftarrow \theta + \delta w$ where δ is the step size. This is called the Natural Actor Critic (NAC) algorithm [68], and has been applied to a number of spoken dialogue systems [28], [69], [70], [71], [72].

NAC is particularly well-suited for on-line operation because the policy is represented as an explicit distribution over actions. Unlike the ϵ -greedy exploration used in Monte-Carlo optimisation (where random actions are chosen with a uniform distribution), the NAC optimisation process can differentiate between less-promising and more-promising actions and therefore explores the solution space more efficiently.

As with LSPI, NAC requires the solution of a system of linear equations (25), which is again cubic in the number

of features, so it is important to choose a compact set of representative features.

F. Section summary

To conclude this section, five approaches to policy representation and optimisation have been presented, which as well as being representative of current practice also illustrate some of the principal design choices. In particular, planning under uncertainty uses the belief state as a probability distribution directly, whereas value iteration and Monte Carlo optimisation require belief space to be quantised, and LSPI and NAC use functional approximations based on weighted linear models of belief state features. Monte Carlo methods and NAC (and SARSA) must be run *on-policy*; planning under uncertainty, value iteration, and LSPI (and Q-Learning) may be run *off-policy*, in *batch*: on-policy methods require that interactions are obtained following the policy under optimization, whereas off-policy methods can perform optimization with interactions obtained by following a different policy. A distinction can also be made between methods like planning under uncertainty and value iteration that perform *planning* on a given model of dynamics, and the other methods described (including Q-Learning and SARSA) which concurrently *learn* about the environment and plan simultaneously. NAC also differs in that it utilises a stochastic policy from which actual actions can be drawn by sampling the policy distribution. This allows a dialogue system to explore alternative actions more effectively by avoiding searching regions of the belief-action space that are very unlikely to ever be visited. Unfortunately, this gain is offset by the generally slow learning characteristics of gradient ascent and in practice the NAC algorithm requires around 10^5 dialogues to optimise a policy [28]. Since the availability of large diverse corpora and/or large numbers of real users willing to interact with a partially trained dialogue system is normally limited, significant attention has been paid to developing user simulators for training statistical dialogue systems and this will be dealt with next.

V. USER SIMULATORS

Learning directly from corpora is problematic since the state space that prevailed during the collection of the data may differ from that used in policy optimisation and also it precludes the use of on-line interactive learning algorithms. An alternative is to build a model of the user that can interact directly with a dialogue system and which can itself be trained on corpora. Such a user simulator can then be used for a wide variety of development, training and evaluation purposes [73].

User simulators normally operate at the abstract level of dialogue acts. Given a sequence of user acts and system responses, the aim is to model the distribution of plausible user responses

$$p(u_t | a_t, u_{t-1}, a_{t-1}, u_{t-2}, \dots) \quad (26)$$

from which an actual user response can be sampled. As indicated in Fig. 2, in a real system the dialogue manager only has access to a noisy observation $o_t = [\tilde{u}_t^1, \dots, \tilde{u}_t^N]$ where \tilde{u}_t^n is a *confused* version of u_t with confidence score

$p(\tilde{u}_t^n | u_t)$. Thus, an error model is needed as well as a user model. The model for $p(u_t | \dots)$ should match the statistics of user responses in available corpora and the error model should match the characteristics of the speech recognition and understanding system [74], [75], [76], [77].

A. User simulation models

One of the earliest approaches to user simulation used N -grams to model (26) directly [78], [79], [80]. The problem with this approach is that large N is needed to capture context and ensure that the user conveys consistent behaviour with the result that the model is inevitably undertrained. A response to this problem was to design simulators that were essentially deterministic and goal-directed but which had trainable random variables wherever there is genuine user choice [81], [82], [83], [84], [85], [86], [87]. These systems can work well but in practice a large amount of hand-crafting is required to achieve acceptable performance.

More recent approaches to user simulation have focused on the use of dynamic Bayesian networks and hidden Markov models paralleling the user model in the POMDP itself [88], [89], [90], [91]. They can also incorporate explicit goals by using a Bayesian network to maintain the goal state of the user [92]. Bayesian network approaches have the advantage that they can model a rich set of conditional dependencies and can be trained on data, although once again data sparsity is a major problem. An alternative promising technique, which avoids the sparsity issues inherent in joint-probability models uses conditional random fields [93]. These have the advantage that they can model very long sequences as features much more efficiently.

Ultimately, the most obvious approach to user simulation will be to train a POMDP-based dialogue system to behave like a user. The simulator could then talk to the system, with the simulator and system each refining their individual policies to maximise reward. The principal barrier to doing this is the lack of an appropriately detailed reward function for the user-side of the dialogue. A solution to this might be to utilise inverse reinforcement learning to infer users' reward functions from real human-human dialogues [94], [95].

B. Error simulation models

User simulation was first introduced to train Markov Decision Process (MDP) based dialogue systems where only the single best output from the recogniser is used. In this case, the error model was required primarily to simulate ASR errors [96], [97], [98], [99], [100]. In contrast, POMDP-based systems use the full distribution of recognised hypotheses at every turn. Hence, as noted at the start of this section, the error model for a POMDP system must generate not just a single confusion but a set of confusions, each with a probability consistent with the behaviour of the actual recogniser. The information content of this distribution makes a significant difference to performance [101]. Hence, it is important to model it accurately.

To conclude this section, user simulation and error modelling are important practical tools for building complex

statistical dialogue systems for which training directly from real users would be impractical. Pragmatically, the approach works reasonably well but the results are highly sensitive to the simulator [102], [103]. POMDP policy optimisation is extremely good at exploiting inconsistent behaviour in a user simulator in order to increase its reward. Thus, it is very easy to obtain very high performance when training and testing on the same simulator, but then to find actual performance in field trials is poor. One way of mitigating against this is to train and test on different simulators [102], [104].

VI. DIALOGUE MODEL PARAMETER OPTIMISATION

Referring back to Fig. 2, a complete POMDP-based dialogue system is characterised by two sets of parameters: the dialogue model \mathcal{M} with parameters τ , which incorporates the user, observation and transition probability distributions; and the policy model \mathcal{P} with parameters θ . Most current POMDP development focusses on policy optimisation and the θ parameters, whereas the dialogue model and the τ parameters are frequently hand-crafted. While hand-crafting the dialog model may seem unsatisfactory, in many cases the dialogue designer will have a strong prior on where the parameters should be. For example, in many applications it may be reasonable to assume that the user goal is constant throughout the dialogue. In this case the probability function $p(s_t|s_{t-1}, a_{t-1})$ becomes a delta function. Similarly, one might reasonably assume that the true user action cannot be a contradiction of the user goal, and the distribution over all actions that are not contradictions is uniform. Hence a user action like “I want a Chinese restaurant” might have zero probability for goals where the food concept is not Chinese but uniform probability otherwise.

In some situations, it is also possible to annotate dialogue corpora to provide complete knowledge of the states of the system. Simple frequency counts can then be used to provide maximum likelihood estimates of key parameters [105]. In [106], maximum likelihood is used to estimate a bigram model of the type of user action while using a handcrafted model to give zero probability to user goals that contradict a given action. In this way simple maximum likelihood can be used to tune relatively complex models.

However, user behaviour and the noise distributions of the speech understanding components are complex and annotating dialogue corpora of sufficient quantity and accuracy is impractical. Hence, key state variables such as the user’s goal must be treated as hidden, and model parameters must be estimated using inference techniques. Recent research is starting to address this problem of learning the dialogue model parameters τ directly from dialogue corpora using a variety of approaches.

A. Expectation maximisation

If the hidden state cannot be annotated then one must turn to an approximate inference algorithm. Since the POMDP has an input-output hidden Markov model structure, the expectation maximisation algorithm (EM) is a sensible first choice. EM operates by using a fixed set of parameters to estimate marginal distributions of the hidden parameters in its first step.

It then uses these marginals to re-estimate the parameters of the model and repeats.

EM has been used to learn the user model for a small spoken dialogue system with 7 states [107]. However, this approach does not scale well when the number of states is significantly increased. The main problem is that marginals must be computed for every possible state, which becomes impractical for many applications. [108] has shown how expectation maximisation (EM) can be used to build models of the user behaviour, although the method requires the user’s goal in the dialogue to remain constant. This is an unsuitable assumption for many real-world dialogues. Nevertheless, approaches based on EM have the major benefit that they do not require annotations of the dialogue state.

B. Expectation propagation

Another algorithm that has been applied to learning the parameters of a spoken dialogue system is expectation propagation [109], [110]. This algorithm operates directly on a factored Bayesian network and simply extends the loopy belief propagation algorithm to handle continuous state spaces. This allows the parameters to be updated during the propagation step and means that all the conditional independence assumptions are used in simplifying the update. This approach has been used to learn the goal evolution parameters of a relatively large tourist information system [111]. The main advantage of this approach is that it requires no annotations of either the true semantics or the dialogue state. This means that parameters can be improved with each dialogue without any extra supervision.

C. Reinforcement learning

All of the above approaches are motivated by the objective of designing a dialogue model that captures the essential elements of the user’s behaviour. An alternative would be to view the dialogue model as just another set of parameters that influence the total reward earned by the system. Hence, the task of the parameter optimisation changes from an inference problem to a reinforcement learning task. Although this means that the parameters are no longer probabilities in the usual sense, the final aim of the system is to maximise its performance and this approach directly optimises this metric.

One algorithm that has been proposed for this purpose is an extension of the natural actor critic (NAC) algorithm discussed in section IV called by analogy “natural belief critic (NBC)” [112]. This algorithm replaces the parameters of the policy that would usually be learned, with the parameters of the user model, denoted τ . This is not entirely straightforward because the reward is not differentiable with respect to τ . However, if the τ parameters are fitted with Dirichlet priors, then the priors are differentiable. Hence, the priors can be optimised using natural gradient ascent, and the required τ parameters can then be obtained by sampling the optimised priors. The NBC algorithm can be further extended to optimise both the policy and the user model parameters at the same time. Experiments have shown that policies optimised by this joint “natural actor and belief critic (NABC)” algorithm

outperform policies trained solely with the natural actor critic algorithm [70].

VII. FAST TRAINING AND USER ADAPTATION

As noted in section IV, policy optimisation for a real-world spoken dialogue system takes an order of $O(10^5)$ training dialogues. This is too large for the policy to be trained in direct interaction with human users and hence, as discussed in section V, policy optimisation currently relies on interaction with user simulators. These simulators require a high level of sophistication and are not trivial to build. Moreover, the policies trained with simulated users are biased towards the particular behaviour that these models incorporate. It is desirable therefore to speed up policy optimisation so that policies can be trained with or adapted to real users.

A second issue with standard policy optimisation techniques is that they do not provide a measure of certainty in the estimates. As dialogue systems become more complex it will be important to convey to the user the system's confidence level. When the system is uncertain, then it needs to signal to the user that it is uncertain. More pragmatically, knowing the level of uncertainty in the various regions of belief-action space will assist learning algorithms to explore parts of the space it is uncertain about instead of exploring randomly.

A final issue concerns the inherent approximation introduced by parameterising the policy. As described in IV, typically a set of feature functions are chosen and then parameters are optimised with respect to the chosen features. To be effective, the feature functions must be carefully hand-crafted and even then the result will be suboptimal.

One method that has been recently applied to spoken dialogue systems to address these issues is Gaussian process-based reinforcement learning, which allows a policy model \mathcal{P} to be defined and optimised in a non-parametric manner⁴. A Gaussian process (GP) is a generative model of Bayesian inference that can be used for function regression [113]. A GP is specified by a mean and a kernel function, which defines prior function correlations and is crucial for obtaining good estimates with just a few observations.

As with the Value Iteration and Monte Carlo methods described in section IV, the GP approach represents the policy via the Q -function (see (9)). However, unlike those methods, it does not require discretisation of belief space. Furthermore, given a suitable kernel, it can operate both in summary belief space and in the full master belief space.⁵

Current research on using GP for dialogue systems is focussed on the GP-Sarsa algorithm [114], which models the Q -function as a zero mean Gaussian process

$$Q(b, a) \sim \mathcal{GP}(0, k((b, a), (b, a))) \quad (27)$$

where $k((b, a), (b, a))$ is a kernel representing the correlations in belief-action space (b, a) .

The posterior of the Gaussian process provides not only the mean estimate of the Q -function but also the variance, which

⁴Here non-parametric does not mean parameter-free but rather that the choice of parameters does not restrict the solution.

⁵Hence in this section, master and summary space are not distinguished.

gives an estimate of the uncertainty in the approximation. In online algorithms, this estimate of the uncertainty can be used for more efficient exploration, either via an active learning model or by defining a stochastic policy model, both of which have been shown to speed up the process of learning [115].

The GP-Sarsa algorithm can be used to estimate the Q -function for real-world problems [116], [115]. It has been shown to optimise dialogue policies faster than a standard reinforcement learning algorithm [117] and it has been used to successfully train a dialogue policy in direct interaction with human users [118].

GP-Sarsa is an on-line on-policy algorithm in that it learns by interaction, taking actions according to the same policy that it is optimising. While on-policy methods guarantee that the overall reward acquired during the optimisation process is maximised, off-policy methods guarantee that a particular policy behaviour is followed during the process of optimisation. Off-policy sample-efficient methods have been explored in the context of parametric optimisation using the framework of Kalman temporal differences [119]. On-line off-policy learning typically uses a form of Q -learning that exploits the Bellman equation

$$Q(b, a) = E \left(r_{t+1} + \gamma \max_{a' \in \mathcal{A}} Q(b_{t+1}, a') \mid b_t = b, a_t = a \right) \quad (28)$$

whose maximisation requires a non-linear parametrisation of the Q -function. The KTD-Q algorithm is a sample efficient algorithm that assumes a non-linear parametrisation of the Q -function and optimises the parameters using Kalman filtering [120]. Similar to GP-Sarsa, this algorithm also provides a measure of the uncertainty of the approximation that can be used for more efficient policy optimisation [120]. This algorithm has been successfully applied to dialogue policy optimisation and shown to perform significantly faster than standard off-policy methods such as LSPI [121], [44].

VIII. SYSTEMS AND APPLICATIONS

The preceding sections of this review provide an outline of the core components of a statistical dialogue system. It will be apparent that these systems are complex and the techniques are being continually refined. Whilst commercial deployment is probably still some way off, several working systems have been built in a number of application areas. This section provides some specific examples of dialogue systems and applications that have been implemented within the POMDP framework.

Most of the systems implemented to date have been information inquiry applications. These include: voice dialling [122], tourist information [42], [28], appointment scheduling [123], and car navigation [24]. Command-and-control applications have also been demonstrated, such as control of devices and services in the home via a multi-modal interface [124]. Another class of application is the troubleshooting domain, which is an interesting example of planning for uncertainty where the observation space of the POMDP can be extended to include non-speech variables such as the state of a modem and the actions can include the invocation of test signals in addition to verbal responses [125].

POMDP-systems have also been demonstrated in the ‘‘Let’s Go’’ challenge run by Carnegie Mellon University, which involves providing live out-of-hours bus information to residents in the Pittsburgh area [34], [126]. This application must deal with casual callers using a variety of mobile phones, often in noisy environments. When success rates were plotted as a function of word error rate, the POMDP-based systems consistently outperformed the conventional baseline system [8].

IX. EVALUATION AND PERFORMANCE

Whilst methods for evaluation of most data-driven tasks in speech and natural language processing are well-established [127], evaluating a spoken dialogue system requires interaction with users and is therefore difficult [128]. Moreover, a spoken dialogue system consists of distinct modules and although there are well-defined evaluation metrics for most of them individually, joint evaluation of the complete system is challenging. Ultimately the goal is usually to provide user satisfaction, but this is hard to measure and is anyway really only appropriate for real users with real needs. Testing a system by giving paid subjects artificial goals will inevitably create biases and even then it is difficult to recruit sufficient subjects to achieve statistical significance for each contrast condition tested. Given these problems, evaluation of SDS typically falls into one of 3 testing regimes: testing with some form of simulated user, testing with paid subjects and testing within a live deployed system.

There are no standardised measures for dialogue manager evaluation. However, the PARADISE framework [129] suggests the general approach of predicting user satisfaction ratings using a corpus of dialogues annotated with user satisfaction ratings and a set of objective measures. It defines the dialogue manager performance as a weighted function of the dialogue success and dialogue-based cost measures such as the dialogue length or the number of times the system produced a confirmation. The weights can be inferred from the annotated corpus using regression, and this regression can then be used as the reward function in reinforcement learning [130], [131]. Whilst this enables a large number of objective measures to be considered, in practice the dialogue success rate and the dialogue length are typically the most important predictors. Hence, most statistical dialogue systems are trained to maximise success whilst minimising the length of the dialogue.

The simplest and most efficient test regime is to use a simulated user. This enables wide coverage of the space of possible dialogues, with a variety of scenarios and the ability to vary the effective recognition error rates over a wide range [132], [133], [134]. An obvious disadvantage, of course, is the potential discrepancy between the behaviour of a simulated user and the behaviour of real users. Nevertheless, user simulation has been a very common way to evaluate different POMDP approaches [125], [28].

As an illustration, Fig. 6 reproduces the results from [28]. Three different dialogue managers were evaluated with a simulated user at varying noise levels. These were a hand-crafted deterministic dialogue manager that takes only the

most likely input from the simulated user (HDC), the BUDS dialogue manager (Section III-B) with a hand-coded policy (BUDS-HDC) and the BUDS dialogue manager with a policy trained using NAC (BUDS-TRA) (Section IV-E). The systems operate in a tourist information domain, where the users may ask about hotels, restaurants, bars and amenities in a fictitious town. The reward function is defined as 20 for a successful dialogue minus the dialogue length⁶. The confusion rate is defined as the probability of a semantic concept from a hypothesis passed to the dialogue manager being incorrect. A dialogue is successful if a suitable venue was offered and all further pieces of information were given. The results in Fig. 6 highlight some important features of the different managers. Firstly, at low error rates all managers perform equally well, but at higher error rates the BUDS systems are clearly more robust to noise. By using belief tracking, the BUDS dialogue manager is better able to handle conflicting evidence from different dialogue turns. Secondly, the BUDS manager with the trained policy gives further improvement compared to the BUDS system with a handcrafted policy, indicating that policy optimisation via reinforcement learning is effective.




Fig. 6. Comparison between the BUDS POMDP-based dialogue manager with a hand-crafted policy (BUDS-HDC) and a trained policy (BUDS-TRA), and a conventional hand-crafted deterministic dialogue manager (HDC) as a function of simulated confusion rate. Each point gives the mean reward for 5000 simulated dialogues. Error bars show one standard error on each side of the mean. Taken from [28].

The second commonly used evaluation regime is to recruit human subjects to test the systems, with the aim of providing more realistic dialogues. Until recently, this has typically been done in a laboratory setting, where paid subjects are given predefined tasks, instructed to talk to the system in a particular way and then asked to rate the dialogues. For example, systems using the POMDP approach have been evaluated in this way and improvements have been demonstrated relative to hand-coded or MDP baselines [49], [135], [28]. Setting aside the issue of how representative paid subjects are of real users, a major problem with this approach is the difficulty and expense of recruiting sufficient subjects. The recent emergence of crowd-sourcing and the ability to transport audio with

⁶This is equivalent to giving a reward of -1 in all non-terminal states and +20 in all terminal states which denote success.

minimal latency over the internet, has greatly mitigated this problem. For example, in [136] the Amazon Mechanical Turk service was used to recruit subjects and provide them with predefined tasks and basic instructions. Subjects then called the dialogue system via a toll-free telephone number and provided feedback after each dialogue. This now provides an effective way of collecting a large number of human-computer dialogues in order to achieve statistically significant results [137]. However, the use of crowd-sourcing does not address the issue of how representative paid subjects are of real users. Indeed, without the ability to monitor subjects in a laboratory setting, the problems of motivation and attention are exacerbated. Furthermore, there is evidence to suggest that crowd-sourced subjects do not rate their dialogue experience accurately [118].

The final evaluation regime is to test systems in the field with real users. This normally requires partnership with a commercial provider or the setting up of a service that the public are naturally motivated to use. Neither of these is easy to arrange and as a response to this problem Carnegie Mellon University (CMU) have set up an evaluation framework called the Spoken Dialogue System Challenge. This utilises a live service for providing bus information for Pittsburgh, Pennsylvania where users are able to call the system and ask for bus information. This service has been provided for some years using human agents during office hours, and by providing a fully automated out-of-hours service, a genuine user demand has been created which by default is serviced by a baseline spoken dialogue system designed by CMU but which can be substituted by other systems for testing. As an example, in the 2010 challenge a number of systems were tested, first by a cohort of students to ensure that the systems were sufficiently robust to put before the public and then they were tested in the live service. Two of those tested were based on the POMDP framework.

The summary results are shown in Fig. 7 from [8], where the logistic regression of dialogue success was computed against word error rate (WER) for each of the systems. It can be seen that the hand-coded baseline system "sys1" is slightly better at low error rates which was to be expected given the long period of development and refinement of that system. However despite the very short development times of the other two POMDP-based systems, both were more robust on high noise levels than the hand-coded baseline system.

X. HISTORICAL PERSPECTIVE

The first reference to the use of POMDPs in spoken dialogue systems is thought to be Roy *et al* in 2000 [9] and Zhang *et al* in 2001 [138]. However, some of the key ideas were being explored before then. The idea of viewing dialogue management as an observable Markov decision process (MDP) and optimising a policy via reinforcement learning is due to Levin and Pieraccini [139], [140]. This work was quickly developed by others [141], [131], [142], [143], [144], [145], [45], [146], [147], [148], [49], [48], [50], [149], [150], [46], [151], [152], [51]. However, the use of MDP approaches lost momentum as it became apparent that the lack of any explicit




Fig. 7. Logistic regression of dialogue success vs word error rate (WER) for the three systems subjected to live testing in the 2010 Spoken Dialogue Challenge. Sys3 and Sys4 were POMDP-based systems and Sys 1 was a hand-coded system. Taken from [8].

model of uncertainty was limiting performance, and although the generalisation to POMDPs was known, their use appeared to be intractable for real-world systems. The first reference to the need for an explicit representation of uncertainty in dialogue is believed to be Pulman in 1996 who proposed modelling dialogue as a *conversational game* [153]. Soon after Heckerman; Horvitz and Paek; and Meng, Wai and Pieraccini showed how Bayesian networks could be used to provide more robust conversational structures [154], [155], [156].

XI. CONCLUSIONS

The development of statistical dialogue systems has been motivated by the need for a data-driven framework that reduces the cost of laboriously hand-crafting complex dialogue managers and which provides robustness against the errors created by speech recognisers operating in noisy environments. By providing an explicit Bayesian model of uncertainty and by providing a reward-driven process for policy optimisation, POMDPs provide such a framework.

However, as will be clear from this review, POMDP-based dialogue systems are complex and involve approximations and trade-offs. Good progress has been made but there is still much to do. There are many challenges, most of which have been touched upon in this review such as finding ways to increase the complexity of the dialogue model whilst maintaining tractable belief tracking; and reducing policy learning times so that systems can be trained directly on real users rather than using simulators. Down the road, there is also the task of packaging this technology to make it widely accessible to non-experts in the industrial community.

There are other challenges too. The POMDP framework depends critically on the notion of rewards. In principle this is a key benefit of the approach since it provides an objective mechanism for specifying dialogue design criteria. However, the problem in practice is that it is very difficult to extract reliable reward signals from users. Even the simple success/fail criterion is difficult to compute since even if asked “Did the system give you all that you asked for?”, many users will say “Yes” regardless just to be polite, or “No” because they had completely unrealistic expectations of what the system can do for them. Reward functions based on user satisfaction as predicted by a regression on objectively measurable features such as in the PARADISE framework [129] may mitigate this problem and need to be explored further [130], [131]. However, much of the experience to date suggests that on-line learning with real users will not be truly effective until robust biometric technology is available that enables the emotional state of the user to be accurately measured [118]. In the meantime, there is no shortage of further areas to develop and explore!

REFERENCES

- [1] R. De Mori, *Spoken Dialogues with Computers*, ser. Signal Processing and its Applications. Academic Press, 1998.
- [2] R. Smith and D. Hipp, *Spoken Natural Language Dialog Systems: A Practical Approach*. Oxford University Press, 1994.
- [3] V. Zue and J. Glass, “Conversational Interfaces: Advances and Challenges,” *Proc IEEE*, vol. 88, no. 8, pp. 1166–1180, 2000.
- [4] S. Young, “Cognitive User Interfaces,” *Signal Processing Magazine*, vol. 27, no. 3, pp. 128–140, 2010.
- [5] M. Oshry, P. Baggia, K. Rehor, M. Young, R. Akolkar, X. Yang, J. Barnett, R. Hosn, R. Auburn, J. Carter, S. McGlashan, M. Bodell, and D. C. Burnett, “Voice extensible markup language (VoiceXML) 3.0,” W3C, W3C Working Draft, Dec. 2009, <http://www.w3.org/TR/2009/WD-voicexml30-20091203/>.
- [6] T. Paek and R. Pieraccini, “Automating Spoken Dialogue Management Design Using Machine Learning: An Industry Perspective,” *Speech communication*, vol. 50, no. 8-9, pp. 716–729, 2008.
- [7] R. P. Lippmann, “Speech recognition by machines and humans,” *Speech communication*, vol. 22, no. 1, pp. 1–15, 1997.
- [8] A. Black, S. Burger, A. Conkie, H. Hastie, S. Keizer, O. Lemon, N. Merigaud, G. Parent, G. Schubiner, B. Thomson, J. Williams, K. Yu, S. Young, and M. Eskenazi, “Spoken dialog challenge 2010: Comparison of live and control test results,” in *Proceedings of SIGDIAL*, Portland, Oregon, 2011, pp. 2–7.
- [9] N. Roy, J. Pineau, and S. Thrun, “Spoken Dialogue Management Using Probabilistic Reasoning,” in *Proceedings of ACL*, Hong Kong, 2000, pp. 93–100.
- [10] S. Young, “Talking to Machines (Statistically Speaking),” in *Proceedings of ICSLP*, Denver, Colorado, 2002.
- [11] J. Williams and S. Young, “Partially Observable Markov Decision Processes for Spoken Dialog Systems,” *Computer Speech and Language*, vol. 21, no. 2, pp. 393–422, 2007.
- [12] L. Kaelbling, M. Littman, and A. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, vol. 101, pp. 99–134, 1998.
- [13] E. Hansen, “Solving POMDPs by searching in policy space,” in *Proceedings of UAI*, Madison, Wisconsin, 1998, pp. 211–219.
- [14] M. L. Littman, R. S. Sutton, and S. Singh, “Predictive representations of state,” in *Proceedings of NIPS*, 2002, pp. 1555–1561.
- [15] R. E. Bellman, *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- [16] J. Pineau, G. Gordon, and S. Thrun, “Point-based value iteration: An anytime algorithm for POMDPs,” in *Proceedings of IJCAI*, Acapulco, Mexico, 2003, pp. 1025–1032.
- [17] T. Smith and R. G. Simmons, “Heuristic search value iteration for POMDPs,” in *Proceedings of UAI*, Banff, Canada, 2004, pp. 520–527.
- [18] M. Spaan and N. Vlassis, “Perseus: randomized point-based value iteration for POMDPs,” *Journal of Artificial Intelligence Research*, vol. 24, pp. 195–220, 2005.
- [19] T. Smith and R. Simmons, “Point-based pomdp algorithms: Improved analysis and implementation,” in *Proceedings of the Twenty-First Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*. Arlington, Virginia: AUAI Press, 2005, pp. 542–549.
- [20] S. Ji, R. Parr, H. Li, X. Liao, and L. Carin, “Point-based policy iteration,” in *Proceedings of AAAI*, Vancouver, Canada, 2007, pp. 1243–1249.
- [21] J. Williams, P. Poupart, and S. Young, “Factored Partially Observable Markov Decision Processes for Dialogue Management,” in *Proceedings of the IJCAI Workshop on Knowledge and Reasoning in Practical Dialog Systems*, Edinburgh, United Kingdom, 2005, pp. 76–82.
- [22] T. Bui, M. Poel, A. Nijholt, and J. Zwiers, “A tractable hybrid DDN-POMDP approach to affective dialogue modeling for probabilistic frame-based dialogue systems,” *Natural Language Engineering*, vol. 15, no. 2, pp. 273–307, 2009.
- [23] S. Young, J. Schatzmann, K. Weilhammer, and H. Ye, “The Hidden Information State Approach to Dialog Management,” in *Proceedings of ICASSP*, Honolulu, Hawaii, 2007, pp. 149–152.
- [24] K. Kim, C. Lee, S. Jung, and G. Lee, “A Frame-Based Probabilistic Framework for Spoken Dialog Management Using Dialog Examples,” in *Proceedings of SIGDIAL*, Columbus, Ohio, 2008, pp. 120–127.
- [25] J. Williams, “Incremental Partition Recombination for Efficient Tracking of Multiple Dialogue States,” in *Proceedings of ICASSP*, Dallas, Texas, 2010, pp. 5382–5385.
- [26] M. Gašić and S. Young, “Effective Handling of Dialogue State in the Hidden Information State POMDP Dialogue Manager,” *ACM Transactions on Speech and Language Processing*, vol. 7, no. 3, 2011.
- [27] B. Thomson, J. Schatzmann, and S. Young, “Bayesian Update of Dialogue State for Robust Dialogue Systems,” in *Proceedings of ICASSP*, Las Vegas, Nevada, 2008, pp. 4937 – 4940.
- [28] B. Thomson and S. Young, “Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems,” *Computer Speech and Language*, vol. 24, no. 4, pp. 562–588, 2010.
- [29] J. Henderson and O. Lemon, “Mixture model POMDPs for efficient handling of uncertainty in dialogue management,” in *Proceedings of ACL*, Columbus, Ohio, 2008, pp. 73–76.
- [30] D. Bohus and A. Rudnicky, “A ‘K hypotheses + other’ belief updating model,” in *Proceedings of the AAAI Workshop on Statistical and Empirical Approaches for Spoken Dialogue Systems*, Boston, Massachusetts, 2006, p. 1318.
- [31] R. Higashinaka, M. Nakano, and K. Aikawa, “Corpus-based discourse understanding in spoken dialogue systems,” in *Proceedings of ACL*, Sapporo, Japan, 2003, pp. 240–247.
- [32] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [33] J. D. Williams, “Using particle filters to track dialogue state,” in *Proceedings of ASRU*, Kyoto, Japan, 2007, pp. 502 – 507.
- [34] B. Thomson, K. Yu, S. Keizer, M. Gasic, F. Jurcicek, F. Mairesse, and S. Young, “Bayesian Dialogue System for the Let’s Go Spoken Dialogue Challenge,” in *Proceedings of SLT*, Berkeley, California, 2010, pp. 460–465.
- [35] P. A. Crook and O. Lemon, “Lossless value directed compression of complex user goal states for statistical spoken dialogue systems,” in *Proceedings of Interspeech*, Florence, Italy, 2011, pp. 1029–1032.
- [36] F. Doshi and N. Roy, “The permutable pomdp: fast solutions to pomdps for preference elicitation,” in *Proceedings of the international joint conference on autonomous agents and multiagent systems*, 2008, pp. 493–500.
- [37] J. Williams and S. Young, “Scaling up POMDPs for dialog management: The ‘summary POMDP’ method,” in *Proceedings of ASRU*, Cancun, Mexico, 2005, pp. 177–182.
- [38] —, “Scaling POMDPs for spoken dialog management,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2116–2129, 2007.
- [39] D. Andre and S. Russell, “State abstraction for programmable reinforcement learning agents,” in *Proceedings of AAAI*, Edmonton, Alberta, 2002, pp. 119–125.
- [40] J. D. Williams, “The best of both worlds: Unifying conventional dialog systems and POMDPs,” in *Proceedings of Interspeech*, Brisbane, Australia, 2008, pp. 1173–1176.
- [41] P. Lison, “Towards Relational POMDPs for Adaptive Dialogue Management,” in *Proceedings of the ACL Student Research Workshop*. Uppsala, Sweden: Association for Computational Linguistics, 2010, pp. 7–12.
- [42] S. Young, M. Gašić, S. Keizer, F. Mairesse, J. Schatzmann, B. Thomson, and K. Yu, “The Hidden Information State Model: a practical

- framework for POMDP-based spoken dialogue management,” *Computer Speech and Language*, vol. 24, no. 2, pp. 150–174, 2010.
- [43] L. Li, J. D. Williams, and S. Balakrishnan, “Reinforcement learning for dialog management using least-squares policy iteration and fast feature selection,” in *Proceedings of Interspeech*, Brighton, United Kingdom, 2009, pp. 2475–2478.
- [44] O. Pietquin, M. Geist, S. Chandramohan, and H. Frezza-Buet, “Sample-efficient batch reinforcement learning for dialogue management optimization,” *ACM Transactions on Speech and Language Processing*, vol. 7, no. 3, pp. 1–21, 2011.
- [45] K. Scheffler and S. Young, “Automatic Learning of Dialogue Strategy using Dialogue Simulation and Reinforcement Learning,” in *Proceedings of HLT*, San Diego, California, 2002, pp. 12–19.
- [46] H. Cuayáhuil, S. Renals, O. Lemon, and H. Shimodaira, “Reinforcement learning of dialogue strategies with hierarchical abstract machines,” in *Proceedings of SLT*, Aruba, 2006, pp. 182–185.
- [47] H. Cuayáhuil, S. Renals, and O. Lemon, “Learning multi-goal dialogue strategies using reinforcement learning with reduced state-action spaces,” in *Proceedings of Interspeech*, Pittsburgh, Pennsylvania, 2006.
- [48] O. Pietquin and T. Dutoit, “A probabilistic framework for dialog simulation and optimal strategy learning,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 2, pp. 589–599, 2006.
- [49] J. Henderson, O. Lemon, and K. Georgila, “Hybrid reinforcement/supervised learning for dialogue policies from Communicator data,” in *Proceedings of the IJCAI Workshop on Knowledge and Reasoning in Practical Dialog Systems*, Edinburgh, United Kingdom, 2005, pp. 68–75.
- [50] M. Frampton and O. Lemon, “Learning more effective dialogue strategies using limited dialogue move features,” in *Proceedings of ACL*, Sydney, Australia, 2006, pp. 185–192.
- [51] J. Henderson, O. Lemon, and K. Georgila, “Hybrid reinforcement / supervised learning of dialogue policies from fixed datasets,” *Computational Linguistics*, vol. 34, no. 4, 2008.
- [52] H. Kurniawati, D. Hsu, and W. Lee, “SARSOP: efficient point-based POMDP planning by approximating optimally reachable belief spaces,” in *Proceedings of Robotics: Science and Systems*, 2008.
- [53] T. H. Bui, M. Poel, A. Nijholt, and J. Zwiwers, “A tractable DDN-POMDP approach to affective dialogue modeling for general probabilistic frame-based dialogue systems,” in *Proceedings of the IJCAI Workshop on Knowledge and Reasoning in Practical Dialog Systems*, Hyderabad, India, 2007, pp. 34–37.
- [54] T. H. Bui, B. W. van Schooten, and D. H. W. Hofs, “Practical dialogue manager development using POMDPs,” in *Proceedings of SIGDIAL*, Antwerp, Belgium, 2007, pp. 215–218.
- [55] F. Pinault, F. Lefevre, and R. D. Mori, “Feature-based summary space for stochastic dialogue modeling with hierarchical semantic frames,” in *Proceedings of Interspeech*, Brighton, United Kingdom, 2009, pp. 284–287.
- [56] A. Boularias, H. R. Chinaei, and B. Chaib-draa, “Learning the reward model of dialogue pomdps from data,” in *NIPS Workshop on Machine Learning for Assistive Techniques*, 2010.
- [57] F. Pinault and F. Lefevre, “Semantic graph clustering for POMDP-based spoken dialog systems,” in *Proceedings of Interspeech*, Florence, Italy, 2011, pp. 1321–1324.
- [58] —, “Unsupervised clustering of probability distributions of semantic frame graphs for pomdp-based spoken dialogue systems with summary space,” in *Proceedings of the IJCAI Workshop on Knowledge and Reasoning in Practical Dialog Systems*, Barcelona, Spain, 2011.
- [59] H. R. Chinaei and B. Chaib-draa, “Learning dialogue pomdp models from data,” in *Canadian Conference on Artificial Intelligence*, 2011, pp. 86–91.
- [60] H. R. Chinaei, B. Chaib-draa, and L. Lamontagne, “Learning observation models for dialogue pomdps,” in *Canadian Conference on Artificial Intelligence*, 2012.
- [61] S. Varges, G. Riccardi, S. Quarteroni, and A. Ivanov, “POMDP concept policies and task structures for hybrid dialog management,” in *Proceedings of ICASSP*, Prague, Czech Republic, 2011, pp. 5592 – 5595.
- [62] J. D. Williams, “Integrating expert knowledge into POMDP optimization for spoken dialog systems,” in *Proceedings of the AAAI Workshop on Advancements in POMDP Solvers*, Chicago, Chicago, Illinois, 2008.
- [63] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, ser. Adaptive Computation and Machine Learning. Cambridge, Massachusetts: MIT Press, 1998.
- [64] F. Lefevre, M. Gasic, F. Jurcicek, S. Keizer, F. Mairesse, B. Thomson, K. Yu, and S. J. Young, “ k -nearest neighbor monte-carlo control algorithm for pomdp-based dialogue systems,” in *Proceedings of SIGDIAL*, London, United Kingdom, 2009.
- [65] M. G. Lagoudakis and R. Parr, “Least-squares policy iteration,” *Journal of Machine Learning Research*, no. 4, pp. 1107–1149, 2003.
- [66] R. J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine Learning*, vol. 8, 1992.
- [67] S. Amari, “Natural gradient works efficiently in learning,” *Neural Computation*, vol. 10, no. 2, pp. 251–276, 1998.
- [68] J. Peters and S. Schaal, “Natural Actor-Critic,” *Neurocomputing*, vol. 71, no. 7-9, pp. 1180–1190, 2008.
- [69] T. Misu, “Dialogue strategy optimization to assist user’s decision for spoken consulting dialogue systems,” in *Proceedings of SLT*, Berkeley, California, 2010, pp. 354 – 359.
- [70] F. Jurčićek, B. Thomson, and S. Young, “Natural Actor and Belief Critic: Reinforcement algorithm for learning parameters of dialogue systems modelled as POMDPs,” *ACM Transactions on Speech and Language Processing*, vol. 7, no. 3, 2011.
- [71] T. Misu, K. Sugiura, K. Ohtake, and C. Hori, “Modeling spoken decision making dialogue and optimization of its dialogue strategy,” in *Proceedings of SIGDIAL*, Portland, Oregon, 2011, pp. 221–224.
- [72] T. Misu, K. Georgila, A. Leuski, and D. Traum, “Reinforcement learning of question-answering dialogue policies for virtual museum guides,” in *Proceedings of SIGDIAL*, Seoul, South Korea, 2012, pp. 84–93.
- [73] J. Schatzmann, K. Weilhammer, M. Stuttle, and S. Young, “A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies,” *The knowledge engineering review*, vol. 21, no. 2, pp. 97–126, Jun. 2006.
- [74] H. Hastie, “Metrics and evaluation of spoken dialogue systems,” in *Data-Driven Methods for Adaptive Spoken Dialogue Systems: Computational Learning for Conversational Interfaces*. Springer, 2012.
- [75] O. Pietquin and H. Hastie, “A survey on metrics for the evaluation of user simulations,” *The knowledge engineering review*, 2011.
- [76] J. Schatzmann, K. Georgila, and S. Young, “Quantitative Evaluation of User Simulation Techniques for Spoken Dialogue Systems,” in *Proceedings of SIGDIAL*, Lisbon, Portugal, 2005, pp. 45–54.
- [77] J. Williams, “Evaluating user simulations with the Cramervon Mises divergence,” *Speech communication*, vol. 50, pp. 829–846, 2008.
- [78] W. Eckert, E. Levin, and R. Pieraccini, “User Modelling for Spoken Dialogue System Evaluation,” in *Proceedings of ASRU*, Santa Barbara, California, 1997, pp. 80–87.
- [79] K. Georgila, J. Henderson, and O. Lemon, “Learning user simulations for information state update dialogue systems,” in *Proceedings of Interspeech*, Lisbon, Portugal, 2005, pp. 893–896.
- [80] —, “User Simulation for Spoken Dialogue Systems: Learning and Evaluation,” in *Proceedings of ICSLP*, Pittsburgh, Pennsylvania, 2006, pp. 45–54.
- [81] K. Scheffler and S. Young, “Probabilistic simulation of human-machine dialogues,” in *Proceedings of ICASSP*, Istanbul, Turkey, 2000, pp. 1217–1220.
- [82] V. Rieser and O. Lemon, “Cluster-based user simulations for learning dialogue strategies,” in *Proceedings of Interspeech*, Pittsburgh, Pennsylvania, 2006.
- [83] O. Pietquin, “Consistent goal-directed user model for realistic man-machine task-oriented spoken dialogue simulation,” in *IEEE International Conference on Multimedia and Expo*, Toronto, Canada, 2006, pp. 425 – 428.
- [84] J. Schatzmann, B. Thomson, K. Weilhammer, H. Ye, and S. Young, “Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System,” in *Proceedings of HLT/NAACL*, Rochester, New York, 2007, pp. 149–152.
- [85] H. Ai and D. J. Litman, “Knowledge consistent user simulations for dialog systems,” in *Proceedings of Interspeech*, Antwerp, Belgium, 2007, pp. 2697–2700.
- [86] J. Schatzmann and S. Young, “The Hidden Agenda User Simulation Model,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 4, pp. 733–747, 2009.
- [87] S. Keizer, M. Gasic, F. Jurcicek, F. Mairesse, B. Thomson, K. Yu, and S. Young, “Parameter estimation for agenda-based user simulation,” in *Proceedings of SIGDIAL*, Tokyo, Japan, 2010, pp. 116–123.
- [88] O. Pietquin, *A Framework for Unsupervised Learning of Dialogue Strategies*, ser. SIMILAR Collection. Presses Universitaires de Louvain, 2004.
- [89] H. Cuayáhuil, S. Renals, O. Lemon, and H. Shimodaira, “Human-Computer Dialogue Simulation using Hidden Markov Models,” in *Proceedings of ASRU*, Puerto Rico, 2005, pp. 290–295.

- [90] H. Cuayáhuitl, S. Renals, O. Lemon, and H. Shimodaira, "Human-computer dialogue simulation using hidden markov models," in *Proceedings of ASRU*, Cancun, Mexico, 2005, pp. 290–295.
- [91] O. Pietquin and T. Dutoit, "A Probabilistic Framework for Dialog Simulation and Optimal Strategy Learning," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 2, pp. 589–599, 2006.
- [92] S. Rossignol, M. Ianotto, and O. Pietquin, "Training a BN-based user model for dialogue simulation with missing data," in *Proceedings of IJCNLP*. Chiang Mai, Thailand: ACLWeb, November 2011, pp. 598–604.
- [93] S. Jung, C. Lee, K. Kim, M. Jeong, and G. G. Lee, "Data-driven user simulation for automated evaluation of spoken dialog systems," *Computer Speech and Language*, vol. 23, no. 4, pp. 479 – 509, 2009.
- [94] A. Ng and S. Russell, "Algorithms for inverse reinforcement learning," in *Proceedings of ICML*. Stanford, California: Morgan Kaufmann, 2000, pp. 663–670.
- [95] S. Chandramohan, M. Geist, F. Lefevre, and O. Pietquin, "User Simulation in Dialogue Systems using Inverse Reinforcement Learning," in *Proceedings of Interspeech*, Florence, Italy, 2011, pp. 1025–1028.
- [96] O. Pietquin and S. Renals, "ASR System Modelling for Automatic Evaluation and Optimisation of Dialogue Systems," in *Proceedings of ICASSP*, Orlando, Florida, 2002, pp. 46–49.
- [97] O. Pietquin and R. Beaufort, "Comparing asr modeling methods for spoken dialogue simulation and optimal strategy learning," in *Proceedings of Interspeech*, Lisbon, Portugal, 2005, pp. 861–864.
- [98] O. Pietquin and T. Dutoit, "Dynamic bayesian networks for nlu simulation with application to dialog optimal strategy learning," in *Proceedings of ICASSP*, Toulouse, France, 2006, pp. 49–52.
- [99] J. Schatzmann, B. Thomson, and S. Young, "Error Simulation for Training Statistical Dialogue Systems," in *Proceedings of ASRU*, Kyoto, Japan, 2007, pp. 526– 531.
- [100] F. García, L. F. Hurtado, D. Griol, M. J. Castro, E. Segarra, and E. Sanchis, "Recognition and understanding simulation for a spoken dialog corpus acquisition," in *Proceedings Text, Speech and Dialogue*, 2007, pp. 574–581.
- [101] B. Thomson, K. Yu, M. Gasic, S. Keizer, F. Mairesse, J. Schatzmann, and S. Young, "Evaluating Semantic-level Confidence Scores with Multiple Hypotheses," in *Proceedings of Interspeech*, Brisbane, Australia, 2008, pp. 1153–1156.
- [102] J. Schatzmann, K. Weilhammer, M. Stuttle, and S. Young, "Effects of the User Model on Simulation-based Learning of Dialogue Strategies," in *Proceedings of ASRU*, Cancun, Mexico, 2005, pp. 220–225.
- [103] D. L. Hua Ai, Joel Tetreault, "Comparing user simulation models for dialog strategy learning," in *Proceedings of NAACL*, Rochester, New York, 2007, pp. 1–4.
- [104] D. Kim, J. Kim, and K.-E. Kim, "Robust Performance Evaluation of POMDP-Based Dialogue Systems," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 1029–1040, 2011.
- [105] J. D. Williams, "Exploiting the ASR N-Best by tracking multiple dialog state hypotheses," in *Proceedings of Interspeech*, Brisbane, Australia, 2008, pp. 191–194.
- [106] S. Keizer, M. Gašić, F. Mairesse, B. Thomson, K. Yu, and S. Young, "Modelling user behaviour in the HIS-POMDP dialogue manager," in *Proceedings of SLT*, Goa, India, 2008, pp. 121–124.
- [107] F. Doshi and N. Roy, "Spoken language interaction with model uncertainty: an adaptive human-robot interaction system," *Connection Science*, vol. 20, no. 4, p. 299318, 2008.
- [108] U. Syed and J. D. Williams, "Using automatically transcribed dialogs to learn user models in a spoken dialog system," *Proceedings of ACL:HLT*, p. 121124, 2008.
- [109] T. Minka, "Expectation Propagation for Approximate Bayesian Inference," in *Proceedings of UAI*. Seattle, Washington: Morgan-Kaufmann, 2001, pp. 362–369.
- [110] B. Thomson, F. Jurčiček, M. Gašić, S. Keizer, F. Mairesse, K. Yu, and S. Young, "Parameter learning for POMDP spoken dialogue models," in *Proceedings of SLT*, Berkeley, California, 2010, pp. 271 – 276.
- [111] B. Thomson, "Statistical methods for spoken dialogue management," Ph.D. dissertation, Cambridge, 2009.
- [112] F. Jurčiček, B. Thomson, S. Keizer, F. Mairesse, M. Gašić, K. Yu, and S. Young, "Natural Belief-Critic: a reinforcement algorithm for parameter estimation in statistical spoken dialogue systems," in *Proceedings of Interspeech*. Tokyo, Japan: ISCA, 2010, pp. 90–93.
- [113] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. Cambridge, Massachusetts: MIT Press, 2006.
- [114] E. Engel, S. Mannor, and R. Meir, "Reinforcement learning with Gaussian processes," in *Proceedings of ICML*, Bonn, Germany, 2005, pp. 201–208.
- [115] M. Gašić, "Statistical dialogue modelling," Ph.D. dissertation, University of Cambridge, 2011.
- [116] Y. Engel, "Algorithms and Representations for Reinforcement Learning," PhD thesis, Hebrew University, 2005.
- [117] M. Gašić, F. Jurčiček, S. Keizer, F. Mairesse, B. Thomson, K. Yu, and S. Young, "Gaussian Processes for Fast Policy Optimisation of a POMDP Dialogue Manager for a Real-world Task," in *Proceedings of SIGDIAL*, Tokyo, Japan, 2010, pp. 201–204.
- [118] M. Gašić, F. Jurčiček, B. Thomson, K. Yu, and S. Young, "On-line policy optimisation of spoken dialogue systems via live interaction with human subjects," in *Proceedings of ASRU*, Hawaii, 2011, pp. 312 – 317.
- [119] M. Geist, O. Pietquin, and G. Fricout, "Kalman temporal differences: The deterministic case," in *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, 2009, pp. 185 –192.
- [120] M. Geist and O. Pietquin, "Managing uncertainty within the KTD framework," in *JMLR: Workshop and Conference Proceedings. Workshop on Active Learning and Experimental Design*, 2011, pp. 157–168.
- [121] O. Pietquin, M. Geist, and S. Chandramohan, "Sample efficient on-line learning of optimal dialogue policies with Kalman temporal differences," in *Proceedings of IJCAI*, Barcelona, Spain, 2011, pp. 1878–1883.
- [122] J. Williams, "Demonstration of a POMDP voice dialer," in *Demonstration Session, Annual Meeting of the ACL*, Columbus, Ohio, 2008.
- [123] S. Janarthnam, H. Hastie, O. Lemon, and X. Liu, "The day after the day after tomorrow?" A machine learning approach to adaptive temporal expression generation: training and evaluation with real users," in *Proceedings of SIGDIAL*, Portland, Oregon, 2011, pp. 142–151.
- [124] K. Kim and G. G. Lee, "Multimodal dialog system using hidden information state dialog manager," in *Proceedings of the international conference on multimodal interfaces demonstration session*, Nagoya, Japan, 2007.
- [125] J. Williams, "Applying POMDPs to Dialog Systems in the Troubleshooting Domain," in *Proceedings of the HLT/NAACL Workshop on Bridging the Gap: Academic and Industrial Research in Dialog Technology*, Rochester, New York, 2007, pp. 1–8.
- [126] —, "An Empirical Evaluation of a Statistical Dialog System in Public Use," in *Proceedings of SIGDIAL*, Portland, Oregon, 2011, pp. 130–141.
- [127] L. Hirschman and H. S. Thompson, *Overview of evaluation in speech and natural language processing*. New York City, New York: Cambridge University Press, 1997, pp. 409–414.
- [128] S. Gandhe and D. Traum, "An evaluation understudy for dialogue coherence models," in *Proceedings of SIGDIAL*, Columbus, Ohio, 2008, pp. 172–181.
- [129] M. Walker, D. Litman, C. Kamm, and A. Abella, "Paradise: A framework for evaluating spoken dialogue agents," in *Proceedings of ACL*, Madrid, Spain, 1997, pp. 271–280.
- [130] E. Levin, R. Pieraccini, and W. Eckert, "Learning dialogue strategies within the Markov decision process framework," in *Proceedings of ASRU*, Santa Barbara, California, 1997.
- [131] S. Singh, M. Kearns, D. Litman, and M. Walker, "Reinforcement learning for spoken dialogue systems," in *Proceedings of NIPS*, 1999.
- [132] T. Watambe, M. Araki, and S. Doshita, "Evaluating Dialogue Strategies under Communication Errors Using Computer-to-Computer Simulation," *IEICE Transactions on Information and Systems*, vol. E81-D, no. 9, pp. 1025–1033, 1998.
- [133] H. Ai and F. Weng, "User simulation as testing for spoken dialog systems," in *Proceedings of SIGDIAL*, Columbus, Ohio, 2008, pp. 164–171.
- [134] J. Schatzmann, "Statistical user and error modelling for spoken dialogue systems," Ph.D. dissertation, University of Cambridge, 2008.
- [135] M. Gašić, F. Lefevre, F. Jurčiček, S. Keizer, F. Mairesse, B. Thomson, K. Yu, and S. Young, "Back-off Action Selection in Summary Space-Based POMDP Dialogue Systems," in *Proceedings of ASRU*, Merano, Italy, 2009, pp. 456–461.
- [136] F. Jurčiček, S. Keizer, M. Gašić, F. Mairesse, B. Thomson, K. Yu, and S. Young, "Real user evaluation of spoken dialogue systems using Amazon Mechanical Turk," in *Proceedings of Interspeech*, Florence, Italy, 2011, pp. 3061–3064.
- [137] I. McGraw, C. Lee, L. Hetherington, S. Seneff, and J. Glass, "Collecting Voices from the Cloud," in *Proc. International Conference on Language Resources and Evaluation*, Malta, 2010, pp. 1576–1583.

- [138] B. Zhang, Q. Cai, J. Mao, E. Chang, and B. Guo, "Spoken Dialogue Management as Planning and Acting under Uncertainty," in *Proceedings of Eurospeech*, Aalborg, Denmark, 2001, pp. 2169–2172.
- [139] E. Levin and R. Pieraccini, "A Stochastic Model Of Computer-Human Interaction For Learning Dialogue Strategies," in *Proceedings of Eurospeech*, Rhodes, Greece, 1997, pp. 1883–1886.
- [140] E. Levin, R. Pieraccini, and W. Eckert, "Using Markov Decision Processes For Learning Dialogue Strategies," in *Proceedings of ICASSP*, Seattle, Washington, 1998, pp. 201–204.
- [141] M. Walker, J. Fromer, and S. Narayanan, "Learning Optimal Dialogue Strategies: A Case Study of a Spoken Dialogue Agent for Email," in *Proceedings of ACL/COLING 98*, Montreal, Canada, 1998, pp. 1345–1351.
- [142] M. Kearns, D. Litman, and M. Walker, "Automatic Detection of Poor Speech Recognition at the Dialogue Level," in *Proceedings of ACL*, College Park, Maryland, 1999, pp. 309–316.
- [143] D. Goddeau and J. Pineau, "Fast reinforcement learning of dialog strategies," in *Proceedings of ICASSP*, Istanbul, Turkey, 2000, pp. 1233–1236.
- [144] S. Young, "Probabilistic Methods in Spoken Dialogue Systems," *Philosophical Transactions of the Royal Society (Series A)*, vol. 358, no. 1769, pp. 1389–1402, 2000.
- [145] S. Singh, D. Litman, M. Kearns, and M. Walker, "Optimizing Dialogue Management with Reinforcement Learning: Experiments with the NJFun System," *Journal of Artificial Intelligence Research*, vol. 16, pp. 105–133, 2002.
- [146] M. Denecke, K. Dohsaka, and M. Nakano, "Learning dialogue policies using state aggregation in reinforcement learning," in *Proceedings of Interspeech*, Jeju Island, South Korea, 2004, pp. 325–328.
- [147] —, "Fast reinforcement learning of dialogue policies using stable function approximation," in *Proceedings of IJCNLP*, 2005, pp. 1–11.
- [148] O. Pietquin, "A probabilistic description of man-machine spoken communication," in *IEEE International Conference on Multimedia and Expo*, Amsterdam, The Netherlands, 2005, pp. 410 – 413.
- [149] J. Tetreault and D. Litman, "Comparing the utility of state features in spoken dialogue using reinforcement learning," in *Proceedings of HLT*, New York City, New York, 2006, pp. 272–279.
- [150] —, "Using reinforcement learning to build a better model of dialogue state," in *Proceedings of EACL*, Trento, Italy, 2006.
- [151] P. Heeman, "Combining reinforcement learning with information-state update rules," in *Proceedings of HLT/NAACL*, Rochester, New York, 2007, pp. 268–275.
- [152] O. Pietquin, "Learning to ground in spoken dialogue systems," in *Proceedings of ICASSP*, Honolulu, Hawaii, 2007, pp. 165–168.
- [153] S. Pulman, "Conversational Games, Belief Revision and Bayesian Networks," in *CLIN VII: 7th Computational Linguistics in the Netherlands meeting*, 1996.
- [154] D. Heckerman and E. Horvitz, "Inferring Informational Goals from Free-text Queries: a Bayesian Approach," in *Proceedings of UAI*, Madison, Wisconsin, 1998, pp. 230–238.
- [155] E. Horvitz and T. Paek, "A Computational Architecture for Conversation," in *Proceedings of the Seventh International Conference on User Modeling*, Banff, Canada, 1999, pp. 201–210.
- [156] H. Meng, C. Wai, and R. Pieraccini, "The use of belief networks for Mixed-Initiative dialog modeling," *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 6, pp. 757–773, Nov. 2003.



Steve Young Steve Young received a BA in Electrical Sciences from Cambridge University in 1973 and a PhD in Speech Processing in 1978. He was elected to the Chair of Information Engineering at Cambridge University in 1994. He was a co-founder and Technical Director of Entropic Ltd from 1995 until 1999 when the company was taken over by Microsoft. After a short period as an Architect at Microsoft, he returned full-time to Cambridge University in January 2001 where he was Head of Information Engineering until 2009 and is now

Senior Pro-Vice-Chancellor.

His research interests include speech recognition, expressive synthesis and spoken dialogue systems. He has written and edited books on software engineering and speech processing, and he has published as author and co-author, more than 200 papers in these areas. He is a Fellow of the UK Royal Academy of Engineering, the IET, IEEE and RSA. He served as the senior editor of *Computer Speech and Language* from 1993 to 2004, and he was Chair of the IEEE Speech and Language Technical Committee from 2008 to 2010. He was the recipient of an IEEE Signal Processing Society Technical Achievement Award in 2004. He was elected a Fellow of the International Speech Communication Association (ISCA) in 2009 and he was the recipient of the ISCA Medal in 2010 for Scientific Achievement.



Milica Gašić Milica Gašić is a Research Associate in the Dialogue Systems Group at the University of Cambridge. She graduated in Computer Science and Mathematics from the University of Belgrade in 2006. Since then she has been at the University of Cambridge first as an MPhil student reading Computer Speech, Text and Internet Technology. Subsequently, she did a PhD in Statistical Dialogue Modelling in 2011. She has published around 20 peer-reviewed conference and journal papers in the area of dialogue management. She served on the

organising committee for Young Researcher's Roundtable on Spoken Dialogue Systems in 2009.



Blaise Thomson Blaise Thomson is a Research Fellow at St John's College in the University of Cambridge. He obtained a Bachelors degree in Pure Mathematics, Computer Science, Statistics and Actuarial Science at the University of Cape Town, South Africa, before completing an MPhil at the University of Cambridge in 2006 and a PhD in Statistical Dialogue Modelling in 2010. He has published around 30 peer-reviewed journal and conference papers, largely in the field of dialogue management. In 2008 he was awarded the IEEE Student

Spoken Language Processing award for a paper at ICASSP and in 2010 he co-authored best papers at SLT and Interspeech. He was co-chair of the 2009 ACL Student Research Workshop and co-presented a tutorial on POMDP dialogue management at Interspeech 2009.



Jason Williams Jason Williams is with Microsoft Research. His interests include spoken dialog systems, planning under uncertainty, spoken language understanding, and speech recognition. He is on the Scientific Committee of SigDial (the Special Interest Group on Dialog and Discourse) and is on the board of directors of the Association for Voice Interaction Design (AVIXD). From 2009-2011, he served on the IEEE Speech and Language Technical Committee (SLTC) in the area of spoken dialogue systems. He holds a PhD and Masters in Speech and Language Processing from Cambridge University (UK), and a BSE in Electrical Engineering from Princeton University (USA). Prior to Microsoft, Jason was Principal Member of Technical Staff at AT&T Labs – Research from 2006-2012. Jason has also held several positions in industry building spoken dialog systems, including at Tellme Networks (now Microsoft) as Voice Application Development Manager.