CogWatch – Cognitive Rehabilitation of Apraxia and Action Disorganisation Syndrome

D3.3.2 Report on Predictive Models: II

Deliverable No.		D3.3.2	
Workpackage No.	**WP3**	Workpackage Title	**Action Recognition and Prediction**
Task No.	**T3.3**	Activity Title	**Action Prediction**
Authors (per company, if more than one company provide it together)		**Author (s) Name (s) (Company)** **Emilie Jean-Baptiste, Martin Russell (UoB), Philipp Gulde, Saskia Steinl, Joachim Hermsdörfer (TUM)**	
Status (F: final; D: draft; RD: revised draft):		**F**	
File Name:		**Cogwatch_D3.3.2_Models_II**	
Project start date and duration		**01 November 2011, 40 Months**	

EXECUTIVE SUMMARY

This report describes the development of predictive models in the CogWatch project. The report is in two parts. The first part, sections 2 and 3, describes the technology that underpins the development of the task model (TM) in P1.2, for the tea-making task, and which will be included in P2 for the tooth-brushing task. The second part, section 4, describes psychological experiments aimed at measuring the extent to which knowledge of body kinematics and gaze can be used to recognise and predict errors in the tea-making task.

The purpose of the TM is to monitor a user's progress through the task, to detect errors and to provide sufficient information to the CogWatch system for useful cues to be created. A number of alternative psychological and mathematical models were considered for the CogWatch TM, and the rationale for choosing a Markov Decision Process (MDP) was presented in deliverable D3.3.1. An MDP-based TM for tea-making has been developed at UOB using a hierarchical description of the task. The TM is implemented in Python and has been fully integrated into the C# environment of the CogWatch system. In addition, a program was written to simulate a user of the system (the 'SimU') based on statistics of sub-goal sequences observed in trials of patients and healthy controls. The SimU allowed a substantial number of 'experiments' to be conducted to evaluate the TM. The evaluations are in terms of the user task completion rate as a function of the accuracy of the action recognition (AR) system, and the compliance of the user (i.e. whether the user follows the cues created in response to the TM outputs). For example, with an AR error rate of 10% the TM achieves a user completion rate greater than 90%.

Some degree of AR error is inevitable. Therefore the most recent research has focussed on developing a TM that is robust against such errors. The new TM uses a Partially Observable MDP. The key difference between the MDP and POMDP TMs is that while the MDP maintains a single estimate of the state of the user in the task, the POMDP 'belief state' is a distribution over all of the MPD states. Experiments have shown that the POMDP-based TM can support user task completion rates of 90% with AR error rates greater than 20%.

Section 4 describes experiments conducted at TUM to measure the relationship between action errors in the tea-making task and user kinematics and gaze. The objective is to provide the psychological basis for incorporating this type of information for error prediction in future TMs. Sixty-seven trials were conducted, of which 41 were performed by controls, 16 by patients with right brain damage (RBD) and 10 by patients with left brain damage (LBD). During trials subjects wore a SMI-ETG eye tracking device and their movements were tracked using the Qualysis system. The data streams were synchronized, segmented into sub-goals, and errors were categorized according to the scheme proposed by Hughes et al. (2013). Analyses were conducted on segments exhibiting errors for error detection, and on the previous segment for error prediction.

The main conclusions of the study are that error recognition and prediction based on monitoring the subjects hand kinematics does not appear to be viable. However, fixation patterns and fixation times do have the potential to support the recognition and prediction of errors. Since gaze leads action, fixations could be used to predict following actions and even errors. Incorporation of eye-tracking and an automatized fixation analysis into the TM could provide a valuable contribution in recognizing and predicting errors.

TABLE OF CONTENTS

1. INTRODUCTION 9

2. ACTION PREDICTION IN THE COGWATCH SYSTEM 10

2.1 Action prediction in the first CogWatch prototype 10

2.2 Outline of the system 11

3. THE COGWATCH TASK MODEL 13

3.1 Description of the MDP-based Task Model 13

3.1.1 Markov decision processes (MDPs) 13

3.1.1.1 Formal definition 13

3.1.2 Interpretation of MDPs in the context of tea-making 13

3.1.2.1 The MDP state space 13

3.1.2.2 The MDP action space 13

3.1.2.3 The MDP transition probabilities 13

3.1.2.4 The Cost Function 14

3.1.2.5 The Optimal Strategy 15

3.1.2.6 The Target Plan 15

3.1.3 Structure of the MDP-based TM in CogWatch 16

3.2 Implementation and Testing 17

3.2.1 The Simulated User (SimU) 17

3.2.1.1 Effects of SimU compliance and PP and N-PP strategies 18

3.2.1.2 Implementation 19

3.3 Coping with uncertainty – the POMDP-based TM 20

3.4 POMDPs 20

3.4.1 Formal definition of a POMDP 20

3.4.2 Action Prediction and Error Recognition in the POMDP-based TM 20

3.4.3 The POMDP belief state space and estimating the optimal strategy 21

3.4.4 Implementation and testing 22

3.5 Summary of the CogWatch Task Models 24

4. ACTION ERRORS IN AADS AND KINEMATICS / GAZE 26

4.1 Task and Procedure 26

4.2 Segmentation 28

4.3 Errors 28

4.4 Relationship between action errors in AADS and kinematics 30

4.4.1 Kinematic measures 30

4.4.2 Results 30

4.4.3 Discussion 32

4.4.3.1 Error classification 32

4.4.3.2 Error recognition 33

4.4.3.3 Error prediction 33

4.4.3.4 Conclusion 33

4.5 Relationship between action errors in AADS and gaze 33

4.5.1 Results 33

4.5.1.1 Performance of healthy control subjects 33

4.5.1.2 Performance of CVA patients 35

4.5.1.3 Gaze and error 38

4.5.2 Conclusion 39

5. CONCLUSIONS 40

TABLE OF FIGURES

Figure 1: Hierarchical tree representation of the "black tea without sugar" version of the tea making task (from D1.1 Report on scenarios)....... 10

Figure 2: The Action Recognition and Prediction (ARP) sub-system of the CogWatch system....... 11

Figure 3: Example of a state transition in the MDP-based Task Model....... 14

Figure 4: Structure of TM in CogWatch....... 17

Figure 5: SimU success rates at varying levels of compliance to the TM's strategy....... 19

Figure 6: Procedure for testing the POMDP-based TM....... 22

Figure 7: Task completion rate on the simple tea-making task as a function of AR error rate for the MDP-based CogWatch Task Model....... 23

Figure 8: Comparison of *SimU* task completion rates on the simple tea-making task as a function of AR error rate for MDP- and POMDP-based Task Models....... 24

Figure 9: Experimental setting for the tea-making task including a water jug, milk, a plate for used teabags, teabags, sugar, coffee, a kettle, a mug and a spoon....... 26

Figure 10: Hand velocity and segments identified for a patients' trial. Note that only segments 1, 2, 3, 4, 6 & 7 were executed and their order was mixed....... 28

Figure 11*:* Average of performed errors per trial for 9 error-classes in controls and the two CVA patient groups. On the right the normalized sum of performed errors per trial....... 29

Figure 12: Mean velocity peak for 4 outlying trials of the misestimation error affected sub-segment 4 versus the corresponding mean of group and condition....... 31

Figure 13*:* Occurrence of execution (EX), misestimation (ME) and quality (Q) errors in the eight sub-segments....... 32

Figure 14: Fixation times in the different sub-segments during the task of tea-making in healthy controls....... 34

Figure 15: Seconds of fixation during the performance of the tea making segments in patients compared to healthy controls....... 36

TABLE OF TABLES

Table 1 - Demographic and clinical data of patients tested in the tea-making task. 27

REVISION HISTORY

Revision no.	Date of Issue	Author(s)	Brief Description of Change
V1	27/10/14	Jean-Baptiste, Russell, Gulde, Steinl, Hermsdörfer	First draft
V2	07/11/14	Hermsdörfer, Gulde	Corrections, minor additions
V3_Final	1/12/14	Russell	Implement corrections in response to reviewer's comments
Final	15/12/14	UPM	Format Review, quality report. Ready for submission.

LIST OF ABBREVIATIONS AND DEFINITIONS

Abbreviation	Abbreviation
AR	**Automatic Action Recognizer**
MDP	**Markov Decision Process**
NPP	**Non Psychologically Plausible**
POMDP	**Partially Observable Markov Decision Process**
PP	**Psychologically Plausible**
SimU	**Simulated User**
TM	**Task Model**

1. INTRODUCTION

The EU CogWatch project is concerned with cognitive rehabilitation of stroke patients who suffer from apraxia and activity disorganisation syndrome (AADS). The purpose of the project is to develop and clinically evaluate interactive technology that can monitor a patient's progress through a task, detect when he or she has made an error, and provide output that can usefully be embodied in a patient cue. From a technology perspective this raises two challenges. The first is automatic action recognition (AR), the ability of the system to automatically recognise the patients' individual actions. The CogWatch AR systems is based on statistical (hidden Markov) modelling of actions, using data from instruments attached to the objects involved in the task. Action recognition in the CogWatch system is described in deliverable D3.2.2. The second challenge is task modelling, which is the ability to use the results of AR to monitor the patient's status in the task, and given that status, to identify and predict errors.

This report describes the development of predictive models in the CogWatch project. The report is in two parts. The first part, Sections 2 and 3, describes the technology that underpins the development of the task model (TM) in P1.2, for the tea-making task, and which will be included in P2 for the tooth-brushing task. The initial TM, based on a Markov Decision Process (MDP) is described and interpreted in terms of the CogWatch task. Experiments are described which use a simulated user to investigate the resilience of the MDP-based TM to AR errors and to non-compliance of the user to cues. The results indicate that task completion rates of 90% can be achieved provided that the AR error rate is no greater than 10%. Motivated by these results, a new TM based on a Partially Observable MDP (POMDP) is described. In a POMDP the user "belief state" is modelled as a probability distribution across all user states, leading to improved robustness against AR error.

The second part of the report, Section 4, describes psychological experiments aimed at measuring the extent to which knowledge of body kinematics and gaze can be used to recognise and predict errors in the tea-making task. Kinematic and gaze data are captured using Qualysis system and an SMI-ETG eye tracking device, respectively. The results suggest that error recognition and prediction based on monitoring the subjects hand kinematics is not viable. However, fixation patterns and fixation times do have the potential to support the recognition and prediction of errors. Fixations could be used to predict following actions and even errors. These results indicate that incorporating eye-tracking and automatized fixation analysis into a future TM could provide a valuable contribution to error recognition and prediction.

2. ACTION PREDICTION IN THE COGWATCH SYSTEM

2.1 Action prediction in the first CogWatch prototype

Action prediction in the first CogWatch prototype is the responsibility of the Task Model (TM). The role of the TM relative to the other components of the Action Recognition and Prediction (ARP) sub-system is shown in figure **Figure 2**.

Recall that the task for the first prototype is tea-making. The tea-making task is described fully in CogWatch deliverable D1.1 “Report on scenarios”. **Figure 1** (taken from D1.1) shows a hierarchical tree based description of one of the tea-making tasks (black tea without sugar). It is included here for completeness.

The root of the tree is identified with the whole task (“prepare a cup of tea”). At the next level this is broken down into “sub-goals” (for example, “heat water”). These sub-goals are broken into “tasks”, which in turn are divided in to “sub-tasks”. In CogWatch the tea-making task is modelled at the sub-goal level.

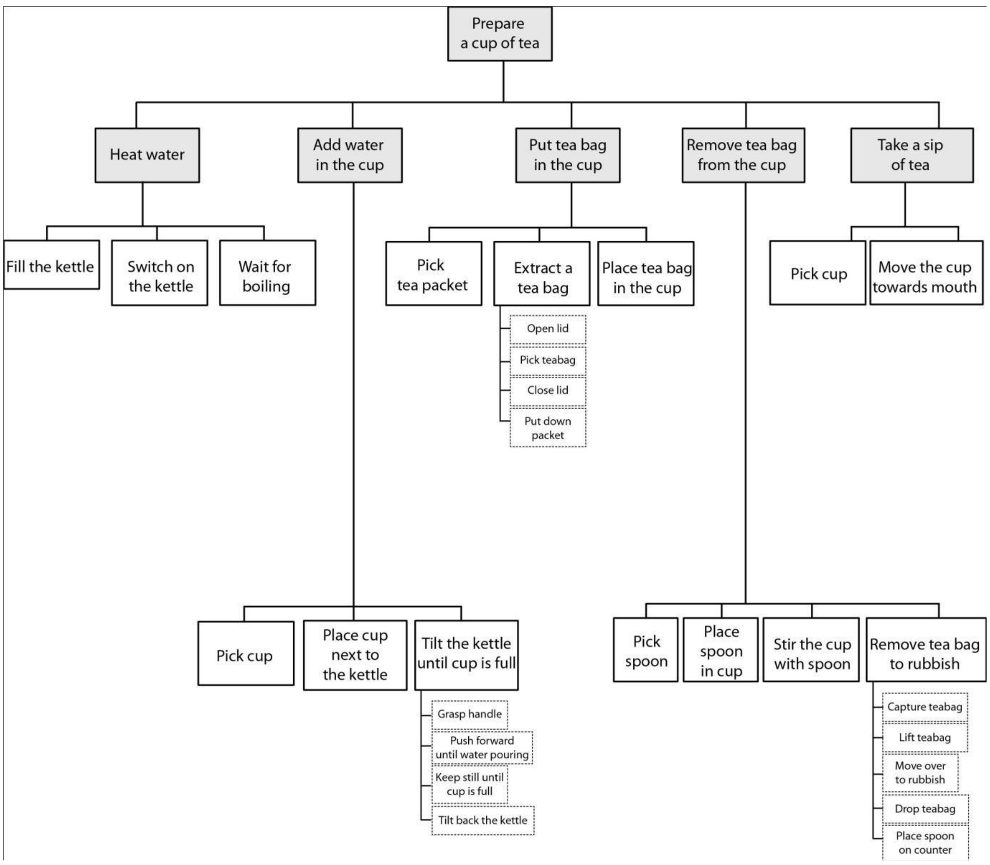

Figure 1: Hierarchical tree representation of the “black tea without sugar” version of the tea making task (from D1.1 Report on scenarios)

2.2 Outline of the system

When the user manipulates the objects involved in a sub-goal of the tea making task, his or her actions cause changes in the outputs of sensors attached to those objects. The sub-goal is recognised from these sensor outputs by the Action Recognition (AR) system.

The AR in the first prototype is described in detail in D3.2.2 “Report on data analysis for action recognition II”. Briefly, the AR is based on statistical models called hidden Markov models (HMMs). The AR system is configured as a parallel set of real-time detectors, one for each sub-goal. This parallel structure was chosen to allow multiple sub-goals to occur simultaneously or at least in overlapping time (which would not be possible in a conventional Viterbi decoder of the type typically used in automatic speech recognition). Each detector contains a sub-gaol model (HMM) and a “toying” model. The role of the toying model is to characterise the range of sensor outputs that occur when the user is not executing the sub-goal. Each detector runs a real-time Viterbi decoder, which decides whether the current input is best described as an instance of the sub-goal or as toying.

The inputs to the TM are sub-goal labels that are output from the AR system. This information is passed to the TM via the VTE.

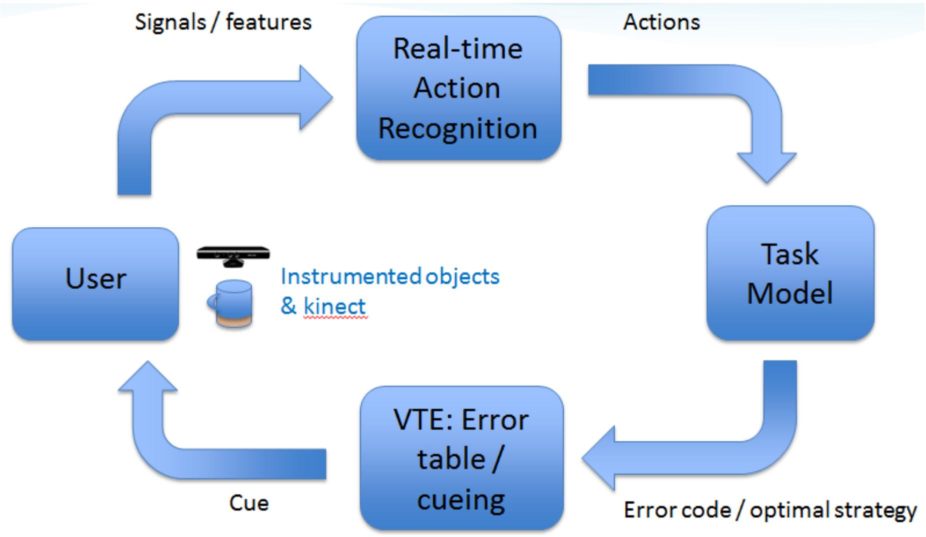

Figure 2: The Action Recognition and Prediction (ARP) sub-system of the CogWatch system.

The purposes of the TM are:

- To estimate the stage that the user has reached in the execution of the task,
- To detect whether or not the user has committed an error, and
- To predict the next sub-goal that the user will execute.

If the TM detects that the user has made an error, then the outputs from the TM are “cue prompts” that are passed to the part of the VTE Information Handler that is responsible for error processing and cueing. If an error has not been detected then the TM sends an indication that the task has been completed successfully. In the event of an error the VTE Information Handler may or may not send a cue to the user, based on the information from the TM.

3. THE COGWATCH TASK MODEL

3.1 Description of the MDP-based Task Model

The TM in CogWatch prototype 1 is based on a Markov Decision Process (MDP). The rationale for choosing an MDP-based TM is explained in D3.2.1.

3.1.1 Markov decision processes (MDPs)

3.1.1.1 Formal definition

Formally, a MDP is a 4-tuple < S , A, $\boldsymbol{P}_a$, $\boldsymbol{R}_a$ >, comprising the following components:

- A finite set $\boldsymbol{S}$ of N states,
- A finite set $\boldsymbol{A}$ of actions,
- For each pair of states $\boldsymbol{s}_1$ and $\boldsymbol{s}_2$ in $\boldsymbol{S}$ and action $\boldsymbol{a}$,
  - $P_a(s_1,s_2)$ is the transition probability of being in state s_2 at time t+1 given state s_1 at time t and that action $\boldsymbol{a}$ was taken
  - $R_a(s_1,s_2)$ is the corresponding reward/cost

3.1.2 Interpretation of MDPs in the context of tea-making

3.1.2.1 The MDP state space

In the context of the tea-making task, a state of the MDP is a sequence of sub-goals that may lead to successful tea-making. Consequently, even for a small number of sub-goals the MDP state space is potentially very large. It is kept finite by restricting the number of repetitions of each sub-goal that are permitted. The size of the state space is further controlled by identifying states that differ only in the number of times that a particular sub-goal has been executed (assuming that it has been executed at least once).

3.1.2.2 The MDP action space

The “actions” component of the MDP consists of the set of sub-goals.

3.1.2.3 The MDP transition probabilities

In the MDP TM, the probability of moving from the current state to another state depends on the current action.

For example, if the current state s_1 corresponds to the sub-goal sequence $a_1,\ldots,a_N$ and the AR outputs sub-goal a, then if $s_2 = a_1,\ldots,a_N, a$ is a valid MDP state (i.e. $a_1,\ldots,a_N, a$ is a valid sequence of sub-goals that can result in successful tea-making) the MDP makes a transition from state s_1 to s_2 with probability 1. If $a_1,\ldots,a_N, a$ is not a valid state then an error has occurred and this information is communicated to the VTE error-handling component of the system.

Figure 3: Example of a state transition in the MDP-based Task Model

3.1.2.4 The Cost Function

The cost function is a mechanism to incorporate human judgment about the importance of different types of behaviour into the MDP. In our case, we combined two types of functions:

i. one that is based on the time taken to complete the task and hence allows the MDP to find the fastest strategy, and

ii. another that takes into account the way participants successfully perform the task and clinicians' preferences. Specifically, we ranked the sub-goals according to the clinicians' priorities, and then associated the highest costs (via the cost function) with failure to implement these most important sub-goals. Consequently, the MDP optimal strategy will prioritize sub-goals that the clinician believes are most important.

Both of these factors are taken into account because the fastest strategy may be valid, but is not necessarily the one that is most psychologically plausible, in the sense that it may not be a strategy that a patient would be likely to pursue or that a clinician would expect. Later we will demonstrate that when the cost function (i) is combined with relevant knowledge from users and clinicians (i.e., cost function (ii)), it allows the TM to generate more meaningful and efficient strategies during the task.

When using the cost function (i) only, the TM's strategies will be referred as Non-Psychologically Plausible; when using the combination (i) and (ii), the TM's strategies will be referred as Psychologically Plausible. Other standard methods will be applied in the future with the aim to compare the TM's performance to other assistive systems.

The cost function includes:

- A penalty based on the time taken to complete the sub-goal.
- A penalty for non-fatal deviations away from the optimal strategy. This is a penalty which is incurred each time the subject executes a sub-goal in a particular state, which is different from the optimal plan or strategy at that state.
- A penalty for repeating a sub-goal (where repetition is not a fatal error). For example if the subject executes a sub-goal “add milk”, then executes one of more other sub-goal, and then executes a second “add milk”, this might incur a penalty even though it is legal.

3.1.2.5 The Optimal Strategy

In the context of MDPs, a *strategy* is a function π: $\boldsymbol{S} \rightarrow \boldsymbol{A}$. In other words, for each state s, $\pi(s)$ is an action (sub-goal).

Given the optimal strategy π^* a function V_π^* can be defined on the state space $\boldsymbol{S}$ such that for any state s, $V_\pi^*(s)$ is the minimum accumulated cost of completing the task given that the participant is currently in state s and follows the optimal strategy π^*.

The cost functions that are used to compute the optimal strategy are based on training material collected during trials and on human intuition. The objective is that the cost function should have the property that large costs are indicative of likely task failure, so that cues can be provided in a timely and psychologically plausible manner. The cost for each sub-goal should be interpreted as the cost for failure to complete that sub-goal. There costs can be based on human intuition, by consulting a clinician to find out which tasks have highest priority, of by analysing training data to find out which sub-goals are given priority in trials.

The optimal strategy is pre-computed using the Monte Carlo Algorithm described in (Levin et al. 2000). It is an iterative algorithm that looks for the action from the Action Space that costs the less to be done after each state contained in the state space. It begins by a guess of what are the best next actions, then improve this guess loops after loops based on the cost functions that are used.

3.1.2.6 The Target Plan

The target plan is a strategy chosen and defined by a human expert, for example a clinician that selects the order in which the sub-goals should be executed, among a specific list of preferred sequences.

This target plan can be selected by hand via the clinician’s screen before the task starts. The information is then passed to the TM for it to retrieve the corresponding strategy for which it would have had been trained beforehand.

3.1.3 Structure of the MDP-based TM in CogWatch

The complete system is shown in Figure 4. In this figure, a_u and a denote the user's action and the TM's strategy, s is the user's state, ER is the error recognition module, and the circumflex indicates an estimate. Error − ID is the user's error type.

The system comprises:

- A set of sensorized objects (specifically a mug, kettle, jug, each fitted with a CogWatch Instrumented Coaster (CIC)),
- A HMM-based AR system,
- A Markov Decision Process based TM, and
- A Prompting System.

The system works as follows: First, the patient chooses the type of tea that he or she wants to make from four options (black tea, black tea with sugar, tea with milk, tea with milk and sugar). This information is passed to the TM and the correct MDP is selected along with the optimal strategy which is computed in advance.

The patient's behavior is detected by monitoring the sensors attached to the objects used during the task. This data is communicated wirelessly to the AR whose aim is to recognise what sub-goal the patient has performed. The AR outputs are passed to the TM, which is in charge of planning and of monitoring the patient's progress through the task.

In other words, each time the patient performs an action (sub-goal), the AAR outputs a sub-goal label, and the TM records it in order to determine the patient's state (i.e, its understanding of what the patient has achieved so far). The state *s* is passed to the Action Policy module that plans what should be done next (for example, what action should be suggested based on the "optimal strategy") in order to assist he patient.

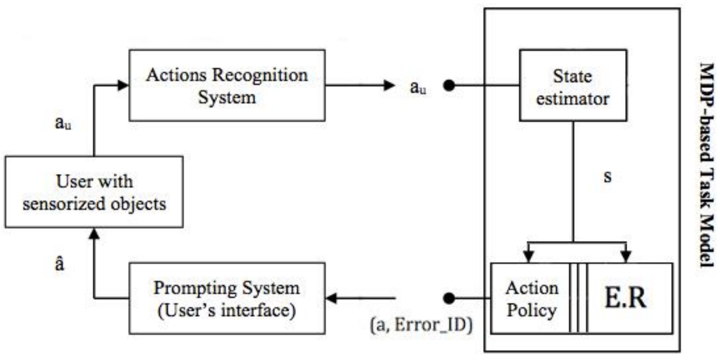

Figure 4: Structure of TM in CogWatch.

In contrast to most previous AI planning systems, we also had to implement an Error Recognition (ER) module. This module analyzes the state *s* in order to identify potential errors in the patient's plan. Finally, the TM outputs a recommendation for the next best action, and if needed alerts the Prompting system that an error has occurred. The Prompting System uses a table designed by clinicians to map the output from the TM to the type of cue that should be retrieved.

Because CogWatch is a rehabilitation system the ER acts as a kind of filter. Each stage of task execution corresponds to a state *s* of the MDP and is associated with an optimal strategy $\pi(s)$. In an assistive technology it might be appropriate to communicate this strategy to the patient at each point in the task. However for a rehabilitation technology this is not the case. The decision about whether or not to convert the optimal strategy into a cue to send to the patient is made by psychologists and coded into the ER system.

3.2 Implementation and Testing

The initial evaluation of the Task Model consisted of measuring its ability to suggest a valid strategy at each stage of task execution. Specifically, the utility of the ER system is not measured here (in fact, the success or otherwise of the ER system can only be measured through repeated trials with patients). In other words, in the present evaluation, success is measured in terms of the TM's ability to suggest an appropriate next action and this is not compromised by the ER system's decision of whether or not to pass it to the patient in the form of a cue.

3.2.1 The Simulated User (SimU)

To reliably evaluate the TM's action policy, a very large number of interactions between participants and the system are necessary. Clearly this cannot be achieved with real human users. Instead, to complete the evaluation, we created a simulated impaired user, *SimU.*

The *SimU*'s behaviour is based on observed sequences of sub-goals from fifty-two control and cognitively impaired participants, aged between 21 and 82, who completed four types of tea making (black tea, black tea with sugar, white tea, white tea with sugar) 100 times. These sequences were used to estimate sub-goal transition probability statistics. Specifically, the statistics that were estimated were:

i. The probability $P_0(a)$ that a is the first sub-goal executed by the patient, and

ii. The transition probabilities $P(a_1 \mid a_2)$ that the patient executes sub-goal a_1 having (immediately) previously completed sub-goal a_2.

The *SimU* generates sequences of sub-goals randomly according to these probabilities. In this way the *SimU* can execute plausible sequences of sub-goals while interacting with a virtualization of the CogWatch system.

A number of experiments were conducted in which the SimU interacted with the simulated CogWatch system. Factors that were considered included:

i. The compliancy of the user. In other words did the simulated user always execute the sub-goal that was suggested by the TM (100% compliancy) or did the simulated user sometimes ignore the suggestions of the TM (less that 100% compliancy)

ii. The psychological plausibility (or otherwise) of the TM's strategy, which in turn derives from the psychological plausibility of the cost function that is used in the computation of the optimal strategy.

The experiments are described below. In summary they show that the MDP-based TM is valid, and that psychologically plausible (PP) strategies are more effective than non-psychologically plausible (NP-P) strategies (Jean-Baptiste et al. 2014).

Figure 5 shows the results of experiments in which the compliance of the SimU was varied between 100% (full compliance, *SimU* always obeys the TM) and 0% (*SimU* ignores the TM), using optimal strategies based on costs functions that ae either psychologically plausible (PP) or not psychologically plausible (N-PP), for each of the four variants of the tea-making task. Performance is measured in terms of the proportion of trials that result in successful completion of the task.

3.2.1.1 Effects of SimU compliance and PP and N-PP strategies

Figure 5 shows The SimU's task completion rate at varying levels of compliance to the Task Model's strategy. PP and N-PP indicate Psychologically Plausible (PP) and Non- Psychologically Plausible (NPP) strategies. The different tea-making tasks are denoted by (a), (b), (d) and (e) and correspond to black tea, black tea with sugar, white tea, and white tea with sugar, respectively.

In Figure 5, we see that if the *SimU* is 100% compliant to the TM's strategies, then whether the TM outputs a NP-P or PP strategy has no impact on the *SimU*'s performance. This indicates that both the NP-P and PP strategies are equally valid, in the sense that if the user follows the TM's instructions the proportion of successful task completions does not depend on whether the cost function used in the computation of the optimal strategy was PP or NPP.

Nevertheless, as soon as the *SimU* decreases its compliance to the TM's outputs, we can see from Figure 5 (a-b-d-e)) that its success rate is higher when the strategies are psychologically plausible than when they are not.

For example, in Figure 5(a-b), when the *SimU* follows a N-PP strategy with a compliance of 20% during the task, its success rate is 78%, which is the same as if it was ignoring the TM (0% compliance). However, under the same circumstances the SimU achieves a task completion rate of 90% with the PP strategy. We can then conclude that if both strategies are valid, the P.P one is optimal compared to N-P.P. To make a parallel with a realistic situation, the P.P strategy can be seen as a familiar one; a strategy able to take into account the ways a clinician would perform the tasks or the optimal ways the patients are used to perform when they succeed. So, with P.P strategies, when the user completes the task and accepts to comply, the TM succeeds to redirect the user on the most efficient ways of succeeding the task (Jean-Baptiste et al., 2014).

On the other hand, even if a N-PP strategy is always correct, it does not take into account the patient's habits, which then leads to more user failures. Indeed, the impact of familiar and unfamiliar sequences on success rate is highlighted in (De Kleine and Van der Lubbe, 2011) and (Graybiel 1998). Familiar sequences are easier to execute and require less effort and energy, as they are controlled through a sub-cortical structure where the sequence is reduced to a single unit. In contrast in the case of novel sequences or sequences that diverge from the familiar, additional cortical mechanisms (more effort, higher demands on resources) are needed.

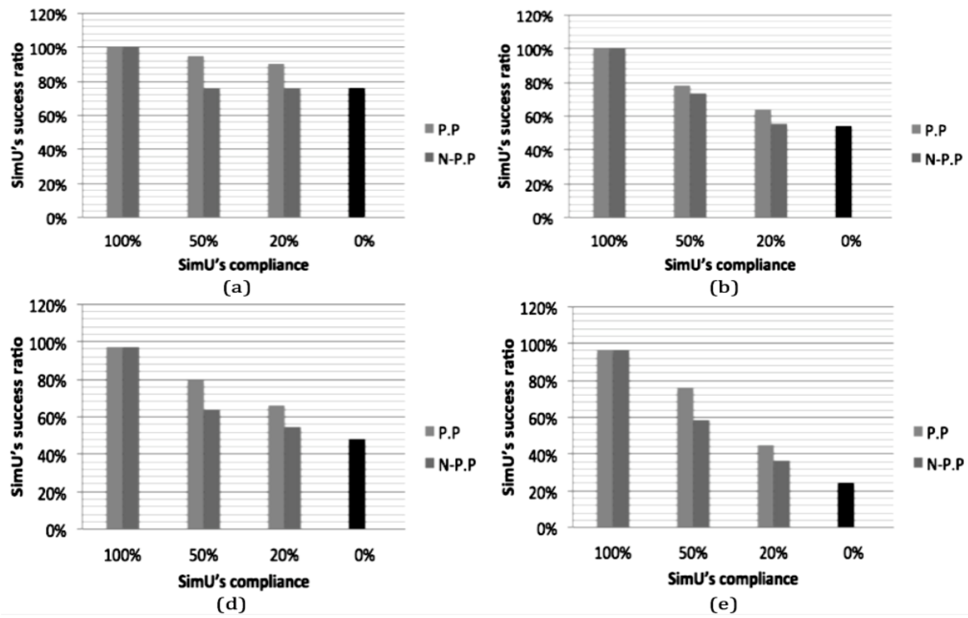

Figure 5: SimU success rates at varying levels of compliance to the TM's strategy.

3.2.1.2 Implementation

The MDP-based TM has been implemented in Python and has been tested successfully via simulation, and with real participants (Pflügler et al. 2014) for tea-making.

Eight patients and three healthy elderly were involved in the experiments.

"In the trials that were performed with the CogWatch system, six out of the nine participants who committed at least one error were successfully cued and achieved the task goal (the selected cup of tea). Two participants did only partly achieve their action goal (e.g. too little of certain ingredient) but were still able to prepare the selected cup of tea. Only one out of the 11 participants was unable to finish his cup of tea due to an irreversible fatal error."

The MDP-based TM has also been implemented and updated in order to be adapted to a new task: teeth-brushing.

3.3 Coping with uncertainty – the POMDP-based TM

A major problem with an MDP-based TM is that it is not well equipped to accommodate errors in the output of the AR system, which are almost inevitable. If the system is in state s_1 then in a conventional MDP the probability $P(s_2 \mid s_1, A)$ of moving to state s_2 depends on s_1 and the sub-goal A that has been executed. In practice, A is unknown and the TM must be satisfied with a, the output the AR system when sub-goal A was performed by the participant.

In this sense the sub-goal is only *partially observable*, since the true sub-goal can only be inferred, and not known, from the recognised sub-goal. Accommodating this uncertainty requires an extension of a MDP called a Partially Observable MDP or POMDP (Williams et al., 2008).

3.4 POMDPs

3.4.1 Formal definition of a POMDP

Formally, a POMDP is a generalization of a MDP. It is a tuple $<\boldsymbol{S}, \boldsymbol{A}, \boldsymbol{P_a}, \boldsymbol{R_a}, \Omega, \boldsymbol{P_o}, \boldsymbol{b_0}>$ where:

- S, A, P_a and R_a define an MDP (section 3.1.1.1).
- Ω is a set of observations.
- $P_o(\boldsymbol{o}, \boldsymbol{a})$ is the probability that an agent will observe $\boldsymbol{o} \in \Omega$ after executing $\boldsymbol{a} \in \mathrm{A}$, reaching state s . The $\boldsymbol{P_o}$ function models the sensor inaccuracy (sensors' noise). In the case of the CogWatch TM this is the error pattern of the AR.
- $\boldsymbol{b_0}$ defines the initial belief state, before the patient has executed an action or received an observation. In our case, the initial state will always be known (i.e., when a new online trial is launched, we consider the state $\boldsymbol{s_0}$ to be empty.).

3.4.2 Action Prediction and Error Recognition in the POMDP-based TM

In the MDP-based TM, the TM has no option other than to "believe" the inputs that it receives from the AR, and bases its action prediction and error recognition on the sequence of sub-goals collected online.

By contrast, the POMDP can accommodate the fact that the AR may misrecognise actions. Instead of basing its optimal strategy on a single state estimate, the POMDP-based TM does this using a probability distribution over all MDP states. The probability distribution over the MDP states is the POMDPs best estimate of where the patient has reached in the task. This distribution is normally referred to as the "belief state".

Each time the user performs a sub-goal, the AR outputs an observation. However, in addition to this observation, the TM also has knowledge of the AR confusion matrix and therefore knows which other actions performed by the user might result in this sub-goal being (erroneously) output and the corresponding probabilities. Using this information the TM replaces the current output sub-goal with a probability distribution over all sub-goals. Combining this distribution with the current belief state creates the new, updated, belief state, which encapsulates the TM's understanding of what the user has achieved so far. Details about exactly how belief states are updated can be found in (Williams et al. 2005).

In summary, the advantage of the POMDP-based TM over the MDP-based TM is that by using knowledge of the typical patterns of error in the AR and representing its belief as a distribution over the MDP states, the POMDP is more resilient to AR error.

Once the belief state is updated, as for the MDP, the POMDP-based TM's action policy module uses adapted cost functions which, taking into account the uncertainty, help defining which next best action should be done by the user, and if an error might have been made.

3.4.3 The POMDP belief state space and estimating the optimal strategy

It was noted in section 3.1.2.1 that even for a simple task such as tea-making the MDP state space is large but discrete and finite. However, because its belief states consist of distributions over the MDP state space, the POMDP state belief space is continuous and infinite. A key property of the MDP TM is the ability to compute the optimal strategy, so that whatever state the patient reaches, the TM can, if needed, suggest the next sub-goal that should be performed in order to minimize the cost of completing the task. For the POMDP to be a useful TM it needs to provide similar information for each of an infinite number of belief states. However, the Monte Carlo algorithm used to obtain the optimal strategy for the MDP-based TM only works for a finite state space.

The solution proposed in (Young et al. 2010) is to use clustering to identify a finite set of belief states that represent the belief state space in some optimal sense. The optimal strategy is then computed for this finite set of belief states using the MDP Monte Carlo algorithm. Finding the optimal strategy for an arbitrary belief state simply involves finding the closest belief state in the finite set and using its optimal strategy. For convenience, the belief states in this finite set will be referred to as "belief centroids".

The process of quantizing the POMDP belief state space involves running the simulated user, *SimU*, many thousands of times in order to populate the belief state space (this will be a small, complex subset of N dimensional space, where N is the number of MDP states). Briefly, each time the simulated user performs a sub-goal a new belief state is created. The distance is computed between this belief state and each of the current belief centroids. If the distance between the current belief state and the closest belief centroid exceeds a threshold then the current belief state becomes a new belief centroid.

From a clustering perspective, this algorithm is unlikely to be optimal and it would be interesting to investigate the application of other clustering algorithms to the creation of the belief centroids.

Once the finite set of belief centroids has been fixed the MDP Monte Carlo algorithm is used to estimate the optimal strategy.

Recent work on the CogWatch project has shown that the choice of the metric (or distance function) used to calculate the distance between a belief state and each of the belief centroids is important. By choosing a suitable metric, the robustness of the POMDP-based TM to AR errors and patient non-compliance to cues can be improved.

3.4.4 Implementation and testing

A POMDP-based TM has been built and tested via simulation for one variation of tea-making. The procedure for testing is illustrated in Figure 6.

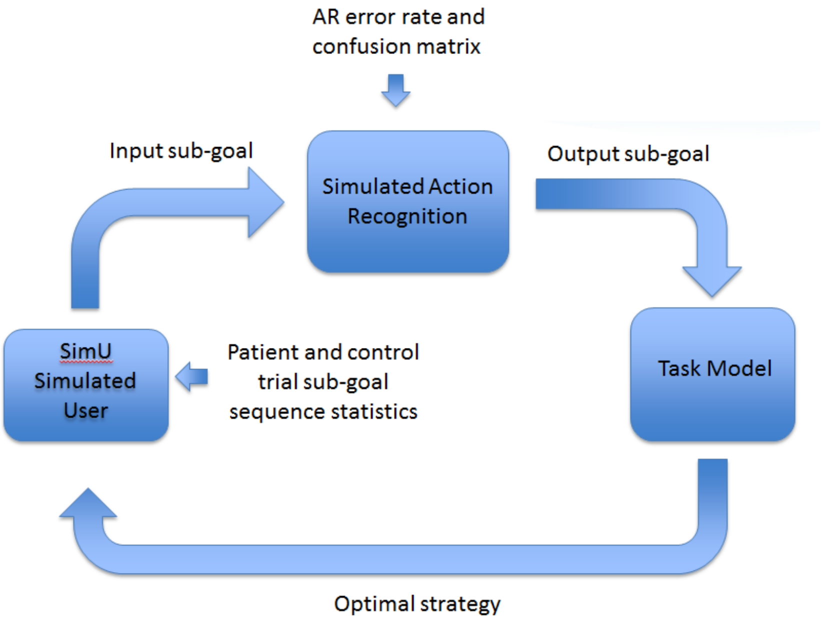

Figure 6: Procedure for testing the POMDP-based TM.

As in the evaluation of the MDP-based TM, a simulated user *SimU* is used to simulate a large number of patient interactions with the system. The *SimU* is a probabilistic model based on the statistics of sub-goal sequences measured in real control and patient trials.

At each step of a simulated trial, the *SimU* generates a sub-goal which is input into the AR. The AR transforms this input sub-goal into the output sub-goal randomly according to the AR error rate and confusion matrix. This is passed to the TM where it, plus knowledge of the AR confusion matrix, is used to update the belief state. The optimal strategy associated with the new belief state is then passed as a cue to the simulated user. If the *SimU* is 100% compliant then it executes the optimal strategy. If the *SimU* is *N%* compliant then it executes the optimal strategy *N*% of the time, and a randomly chosen strategy (100-*N*)% of the time.

From AR error rates of between 0% and 60%, we evaluated the impact of the MDP-based TM's strategies on the *SimU*'s success rate, in the case where the latter was 100% compliant to the TM. From Figure 7, one can see that when the AR's error rate increases, the *SimU*'s ability to complete the task successfully decreases.

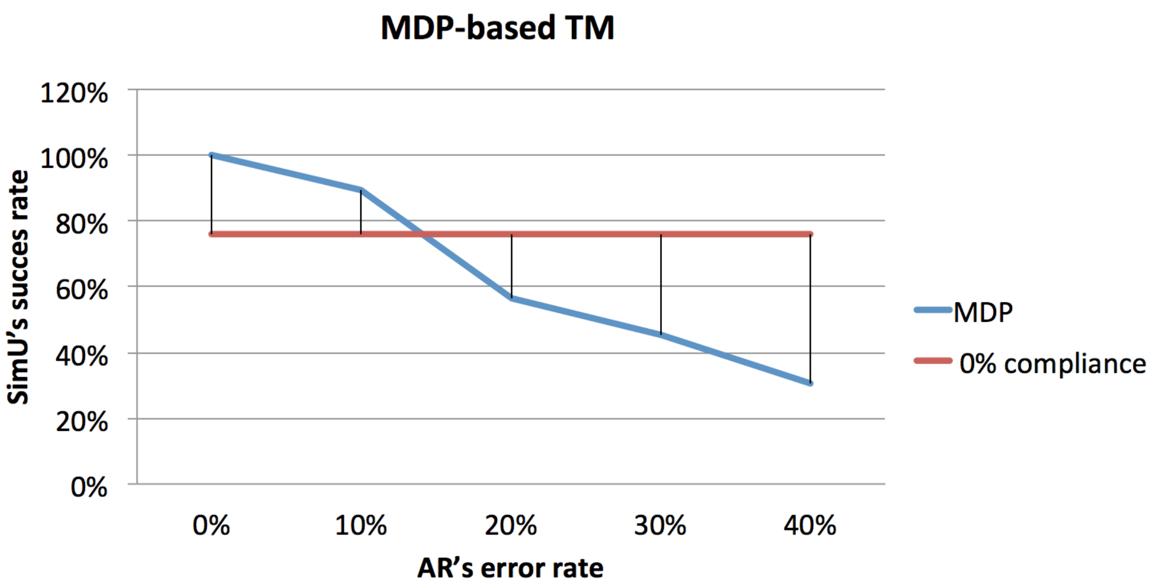

Figure 7: Task completion rate on the simple tea-making task as a function of AR error rate for the MDP-based CogWatch Task Model

This is due to the fact that the MDP-based TM cannot cope with uncertainties related to the AR's outputs. With the MDP, the TM believes that the AR outputs perfectly correspond to the *SimU*'s actions: it does not take into account potential AR action recognition errors. So, as we increase the AR error rate, we increase its potential to misrecognize the action performed by the *SimU.* This increases the probability that the TM has an incorrect understanding of the *SimU*'s history of actions, which then increases the probability that it will deliver an incorrect cue to the *SimU*, who, being 100% compliant, will take the wrong action. Consequently the *SimU* task failure rate will increase.

With a POMDP-based TM, the system acknowledges the fact that the AR can make mistakes. Instead of considering that the AR's output directly correspond to the *SimU*'s action, the POMDP-based TM treats it as an observation from a random variable and updates its belief state (i.e. its understanding of what the user has achieve so far, which is encoded as a probabilistic distribution over the MDP states) according to its understanding of the distribution of SimU actions that could give rise to the observed AR output.

Figure 8 shows a comparison of the task completion rates achieved by the SimU with MDP- and POMDP-based Task Models for the simple tea making task. In this case the metric used in the quantisation of the POMDP belief state space was correlation distance. For example, it is evident that the task completion rate with an AR error rate of 20% and using a POMDP-based TM is better than that achieved with the MDP-based TM and an AR error rate of just 10%. At 20% AR error rate, the task completion rate for the POMDP-based TM is over 90%, compared with less than 60% for the MDP-based TM.

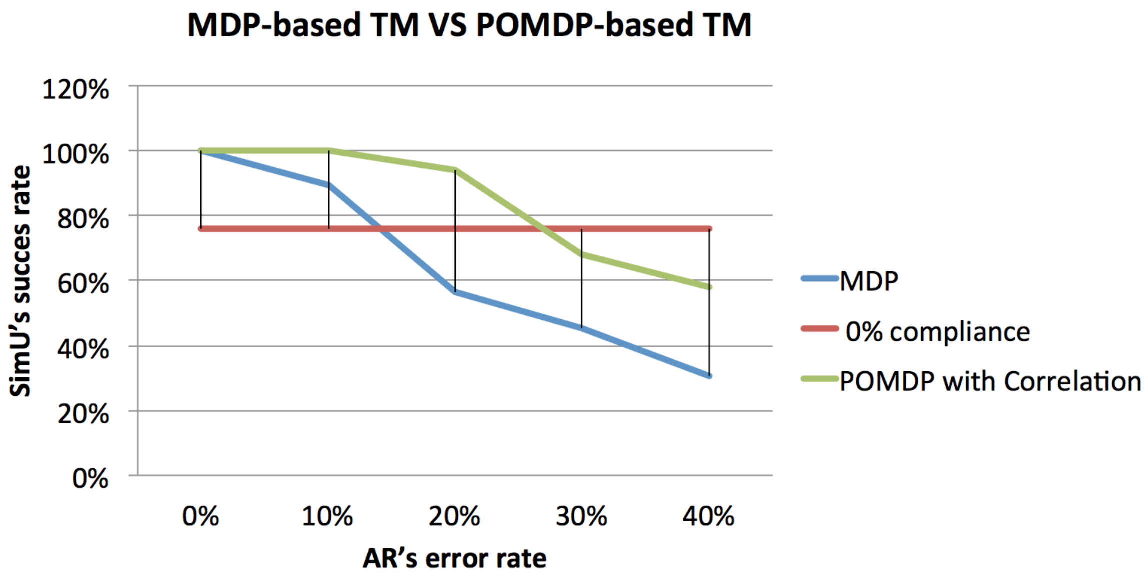

Figure 8: Comparison of *SimU* task completion rates on the simple tea-making task as a function of AR error rate for MDP- and POMDP-based Task Models.

In Figure 7 and Figure 8, the line labelled "0% compliance" corresponds to what the *SimU's* success rate is when it performs the trials by itself (i.e. without any help from the TM). We can see that at around 13% AR error rate, it is better for the *SimU* to ignore the cues based on the optimal strategy of the MDP-based TM and instead to perform the task by itself, according to its own internal model. Thus, we can say that at around 13% AR error rate, the MDP-based TM has no practical utility. On the other hand, for the POMDP-based TM the same phenomenon does not occur until the AR error rate reaches around 28%.

We believe that these results have significant practical implications. In our current AR experiments (reported in deliverable D.3.2.2) the AR error rate is typically below 10% and therefore within the range that can be accommodated by the current MDP-based TM. However, if AR performance degrades when the system is used by patients, or because cheaper or fewer sensors are used to instrument the task objects, then it may be possible to accommodate the degradation in AR performance by upgrading the MDP-based TM to a POMDP-based TM. Similarly, if the system is applied to a new and more demanding task, any increase in AR error rate is more likely to be accommodated with a POMDP-based TM.

3.5 Summary of the CogWatch Task Models

Section 3 has described an approach to task modelling based on Markov Decision Processes (MDPs), which has been used in the first CogWatch prototype for the tea-making tasks.

Our experiments have shown that the MDP-based TM integrated into a simulation of the current CogWatch System can correctly assist a virtually impaired simulated user to complete the task, provided that the AR error rate is not too high. We believe that similar results will be observed when experiments being run with real participants are completed.

From an architectural and computational point of view, implementing this virtual simulation of CogWatch allowed us to validate the TM's capability to fulfill the requirements needed for the system to be a context-aware, intelligent, assistive device.

Although the first prototype focused on the tea-making task, the TM's structure flexibility has allowed it to be extended to the P2 teeth-brushing task.

We saw that a limitation of the MDP-based approach is that it is not well-suited to coping with ambiguity in its inputs. In this application ambiguity arises as a consequence of classification errors in the AR. A Partially Observable MDPs (POMDP) was proposed as a solution to this problem. Results obtained with the POMDP-based TM show how it clearly outperforms the MDP-based TM at higher AR error rates.

4. ACTION ERRORS IN AADS AND KINEMATICS / GAZE

The previous sections describe progress on the "technological" aspects of CogWatch in the area of error prediction through the development of different statistical task models. The following sections describe progress on understanding how errors in AADS can be predicted using information about the kinematics and gaze of patients and healthy subjects.

4.1 Task and Procedure

In the version of the tea-making task tested here, participants are instructed to prepare a cup of tea with milk and one sugar cube. Thus, the following items were placed on the table: a kettle, teabags, milk, sugar cubes and an additional distractor item (instant coffee jar).

The following conditions were tested:

- Bimanual: use of both hands
- Unimanual: use of the ipsilesional hand in patients and the dominant hand in healthy subjects, respectively
- Unimanual: use of the contralesional hand in patients and the non-dominant hand in healthy subjects, respectively

Every condition was repeated once, resulting in a maximum of six trials. The order was bimanual, unimanual (ipsilesional / dominant), unimanual (contralesional / non-dominant).

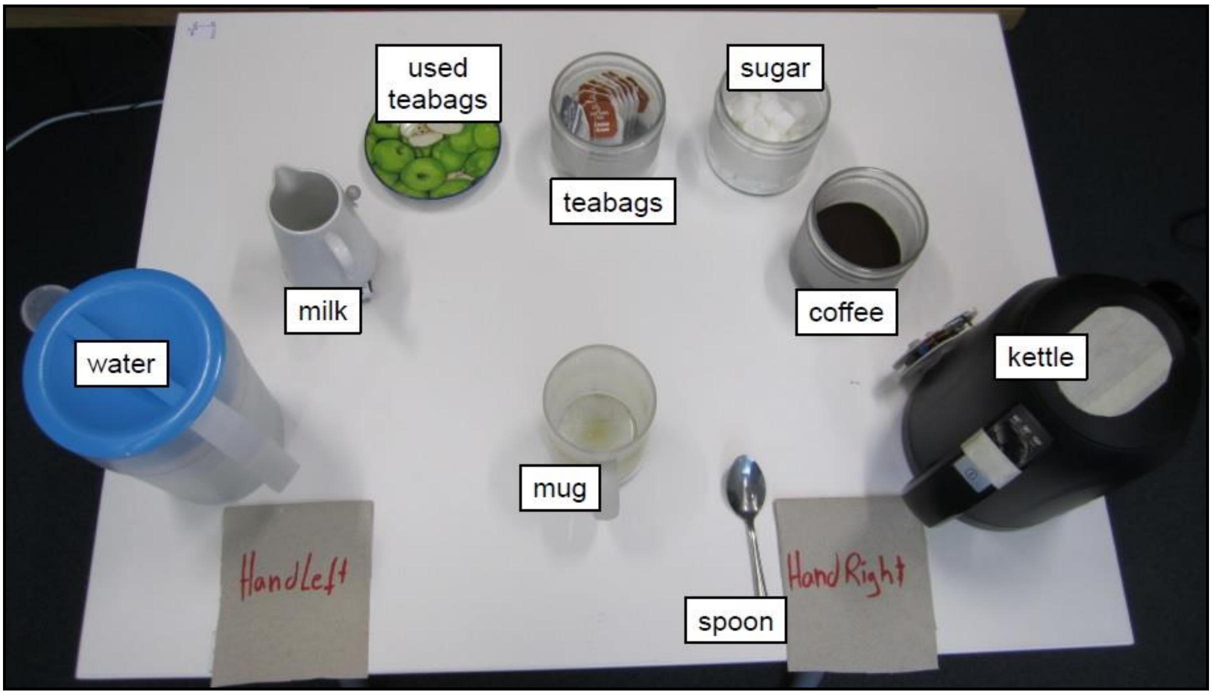

Figure 9: Experimental setting for the tea-making task including a water jug, milk, a plate for used teabags, teabags, sugar, coffee, a kettle, a mug and a spoon.

The settings of the objects available in the task are shown in Figure 9. Starting positions for the left and the right hand are represented by the labeled papers. In the beginning of each trial, the water jug is filled with approximately 0.5 liters of preheated water, the milk carafe is filled, the teabag labels are prepared for an easy entanglement, particular in unimanual trials, and the kettle body is empty. The containers' handles are directed towards the subject.

Until now 9 controls and 7 CVA patients (3 with left brain and 4 with right brain damage) were tested and analyzed. There were 67 trials in total, of which 41 were performed by controls, 16 by patients with right brain damage (RBD) and 10 by patients with left brain damage (LBD). Patients were recruited from the Clinic for Neuropsychology at the Hospital München-Bogenhausen in Munich. Patient's age ranges from 47 to 79 years with a mean of 63 (±9.11) years and time since stroke between 0.5 and 6.5 years with a mean of 2.5 (±2.2) years. Controls had a mean age of 70.88 (±3.4) years. One of the LBD patients, 4 of the RBD patients and 3 of the control subjects were male. Subjects were tested for handedness by the Edinburgh Handedness Inventory. All CVA patients but one LBD patient, and all controls subjects but one were right handers, almost all of them strong (13).

Table 1 - Demographic and clinical data of patients tested in the tea-making task.

Code	Age	Sex	Side of Brain Damage	Paresis	Time since Stroke	EHI
S20	47	M	Left	Yes	1y	1
S22	70	W	Left	Yes	0.5y	1
S36	63	W	Left	Yes	0.5y	0
S85	70	M	Right	Yes	6.5y	0.68
S93	58	W	Right	Yes	3y	1
S96	79	M	Right	Yes	4y	1
S115	64	M	Right	Yes	2y	0.8
Mean	*63 ± 8.33*				*2.5 ± 2.2*	

Subjects are asked to wear a SMI-ETG eye tracking device during task performance. The eye tracking glasses incorporate a HD scene camera with a sampling rate of 30Hz. Fixations were identified and assigned to fixated objects off-line.

Positional data of both hands were recorded with the use of 5 Oqus Motion Capture cameras included in a Qualysis motion capturing system with a sampling frequency of 120Hz. Three passive markers were attached to each hand in the middle of the dorsum. For the analysis only one marker was used, the other two were attached for a better recording reliability and in case one or two markers got lost.

The mug, the milk carafe and the kettle's base and body had force and acceleration sensors with a sampling rate of 200Hz (now 50Hz, see elsewhere in this report) attached. These instrumented coasters are custom made by the University of Birmingham, UK.

4.2 Segmentation

In a first step data streams were synchronized using MatLab and a Visual C executable file. Then data were segmented into discrete actions and analyzed. The coarse boundaries of the action segments were manually defined via the SMI-ETG HD scene camera's video data. The fine adjustment of action-segments is then performed with the use of hand kinematics in MatLab. The whole task is segmented into the following eight action segments (Humphreys & Forde, 1998 in Forde et al., 2010):

1. pour water in the kettle
2. switch the kettle on
3. place a teabag in the mug
4. pour heated water into the mug
5. remove the teabag
6. add milk
7. add one sugar cube
8. stir the tea

Figure 10 shows an example of action segmentation for one trial performed by a patient.

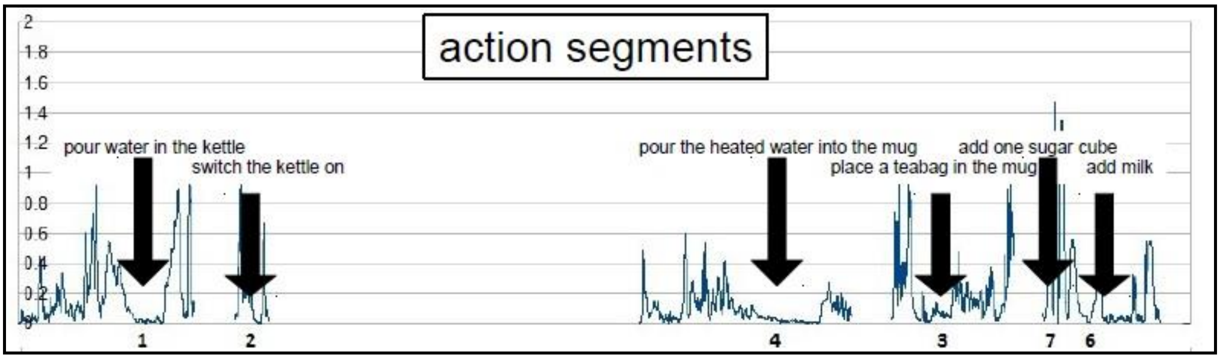

Figure 10: Hand velocity and segments identified for a patients' trial. Note that only segments 1, 2, 3, 4, 6 & 7 were executed and their order was not strictly ascending.

4.3 Errors

For the analysis of errors performed in the task, the error classification of Hughes et al. (2013) was applied. This classification uses 12 different kinds of errors occurring in the tea-making task:

• Addition (AD)	*adding an extra component action that is not required in the action sequence*
• Anticipation (AN)	*performing an action earlier than usual*
• Execution (EX)	*an error in the execution of the task*

- Ingredient omission (IO) *failing to add an ingredient required to complete the task goal*
- Misestimation (ME) *using grossly too much or too little of some substance*
- Mislocation (ML) *an action that is appropriate to the object in hand but is performed in completely the wrong place*
- Ingredient substitution (IS) *an intended action is carried out but with an unintended ingredient*
- Perseveration (PER) *the unintentional repetition of a step or subtask*
- Object substitution (OS) *an intended action carried out with an unintended object*
- Quality (Q) *the action was carried out, but not in an appropriate way*
- Sequence (S) *performing an action much later than usual*
- Sequence omission (SO) *an action sequence in which one step or subtask is not performed, despite the lack of any intention to omit the step or subtask*

Of these errors, 9 have been more or less frequently observed in the subjects performing the tea-making task. Figure 11 shows the average of errors per trial for each of these 9 error-classes and the normalized (control = 1 with SD of 0.15) sum of errors per trial for the three groups of LBD and RBD patients and controls.

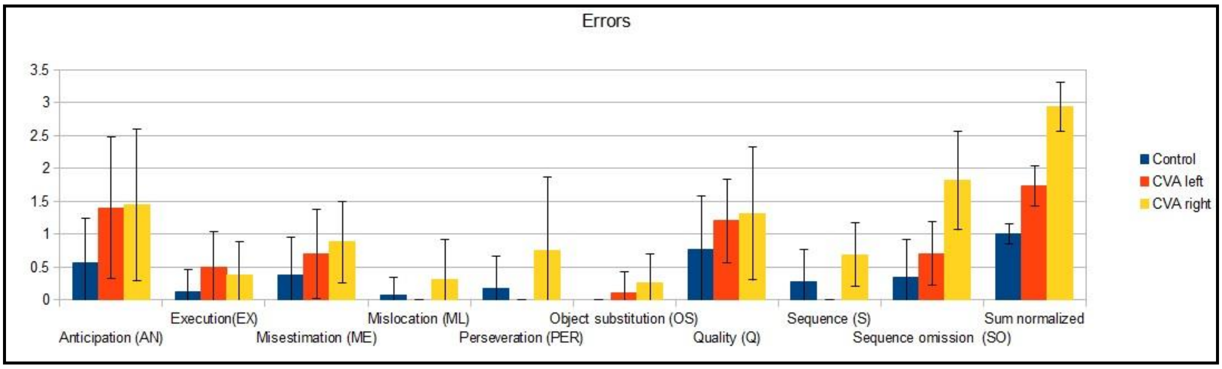

Figure 11*:* Average number of errors per trial for 9 error-classes in controls and the two CVA patient groups. On the right the normalized sum of errors per trial is indicated.

Note also that the control subjects performed a few errors but the patient group committed a higher number of errors, especially anticipation, execution, misestimation, quality and sequence omission errors. Interestingly the frequency of perseveration, sequence and sequence omission errors differed between patients with right and patients with left-sided brain damage but not anticipation errors. This indicates that right sided lesions more frequently cause problems organizing action sequences (ADS). Overall the LBD patients of the sample performed fewer errors than the RBD patients.

4.4 Relationship between action errors in AADS and kinematics

4.4.1 Kinematic measures

The positions and velocities of the hands were determined from the motion recordings and smoothed using a 1s LOESS filter ('local regression'). The following measures were determined for the complete action and the action segments:

- maximum peak velocity
- mean velocity peak
- number of velocity peaks
- movement times
- path lengths

The *'maximum peak velocity'* describes the maximum tangential speed reached in the segment respectively in the full trial.

The *'mean velocity peak'* describes the average local maxima of the speed, so the subject's pace in the discrete movements ignoring pauses (not to be confused with the average speed).

The *'number of velocity peaks'* is calculated as an indicator of movement smoothness, since the tea-making task is a composition of mostly discrete movements. Caution is requested when comparing the differences between sub-segments, since the number of peaks in a reaching movement and in stirring the tea express different aspects of smoothness.

The *'movement time'* is the time taken to complete the single sub-segments respectively the whole task without the waiting period for the boiling of the water which is usually distinguished by resting hands.

The *'path length'* is the tangential distance traveled by the left and the right hand. The path length can be increased due to additional action as well as less goal-directed movements, changes of directions or even tremor.

4.4.2 Results

An analysis of kinematic peculiarities in sub-segments with error occurrence has so far been done for the errors ‘execution’, ‘misestimation’ and ‘quality’. The affected sub-segment was analyzed for error recognition and the preceding sub-segment for error prediction. Measures were compared with the mean of the sub-segment, group and condition with ± 2 standard deviations.

Execution (EX)

LBD patients performed 5 execution errors in 10 trials, RBD patients 6 in 16 trials and controls 5 in 41 trials. Affected sub-segments were:

- #1 ‘pour water in the kettle’ in 4 cases
- #3 ‘ place a teabag in the mug’ in 3 cases

- #4 ‘pour the heated water in the mug’ in 4 cases
- #6 ‘add milk’ in 3 cases
- #8 ‘stir the tea’ in 2 cases

The analyses of the affected and of the preceding sub-segments showed only three outliers in 145 comparisons with no logical structure.

Misestimation (ME)

LBD patients performed 7 misestimation errors in 10 trials, RBD patients 14 in 16 trials and controls 15 in 41 trials. The affected sub-segments were:

- #1 ‘pour water in the kettle’ in 6 cases
- #4 ‘pour the heated water in the mug’ in 24 cases
- #6 ‘add milk’ in 6 cases

Note that only pouring actions were affected by misestimation errors in the tea-making task. The analyses for the affected and for the preceding sub-segments showed 24 outliers in 390 comparisons. In sub-segment 4, affected by misestimation errors, subjects showed in 5 cases (4 trials, one of them bimanual) increased mean velocity peaks.

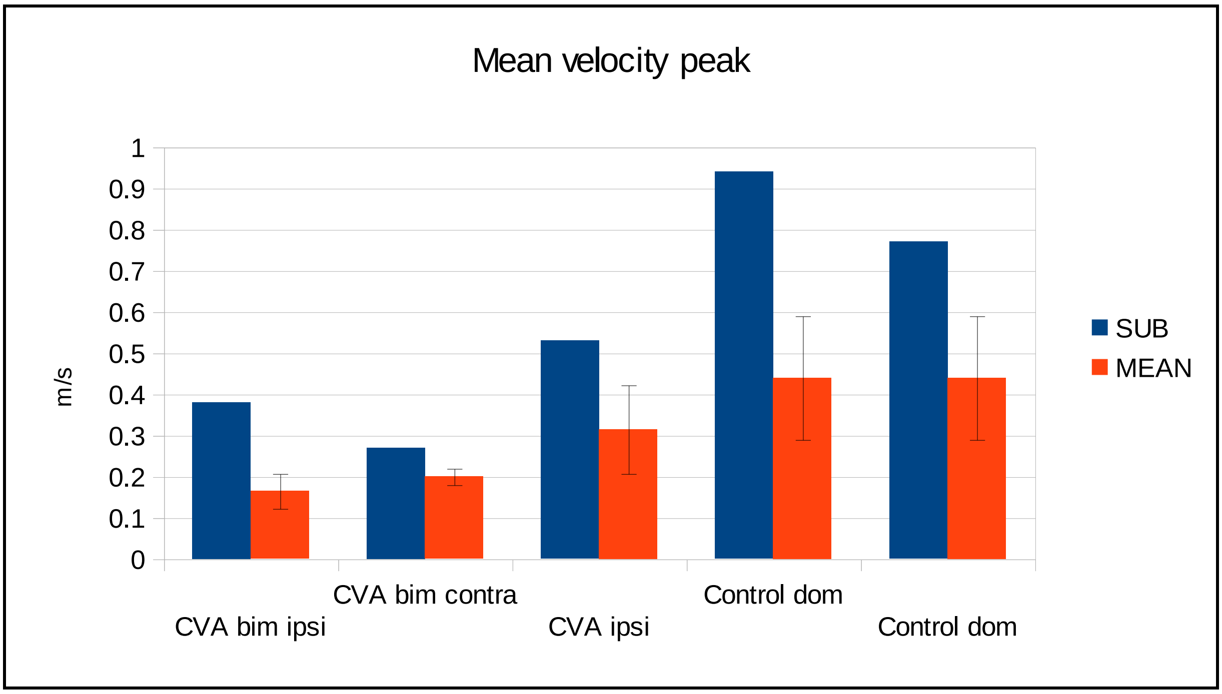

Figure 12: Mean velocity peak for 4 outlying trials (bimanual values are from 1 trial) affected by misestimation errors in sub-segment 4 versus the corresponding mean of group and condition.

Although only in 4 out of 24 trials, subjects showed that this finding might be a first approach for error recognition via kinematics. The other measures and comparisons did not reveal any consistent pattern of kinematic features corresponding with misestimation errors in the sub-segments directly affected or preceding the affected sub-segment.

Quality (Q)

LBD patients performed 12 quality errors in 10 trials, RBD patients 21 in 16 trials and controls 31 in 41 trials. Affected sub-segments were:

- #1 ‘pour water in the kettle’ in 35 cases
- #3 ‘ place a teabag in the mug’ in 6 cases
- #4 ‘pour the heated water in the mug’ in 10 cases
- #5 ‘remove the teabag’ in 1 case
- #6 ‘add milk’ in 2 cases
- #7 ‘add one sugar cube’ in 2 cases
- #8 ‘stir the tea’ in 4 cases

The analyses for the affected and for the preceding sub-segments revealed only 19 outliers in 610 comparisons, with no logical pattern. Due to a quite complex mechanics for opening the kettle’s lid, segment 1 resulted in many quality errors in controls as well as in CVA patients.

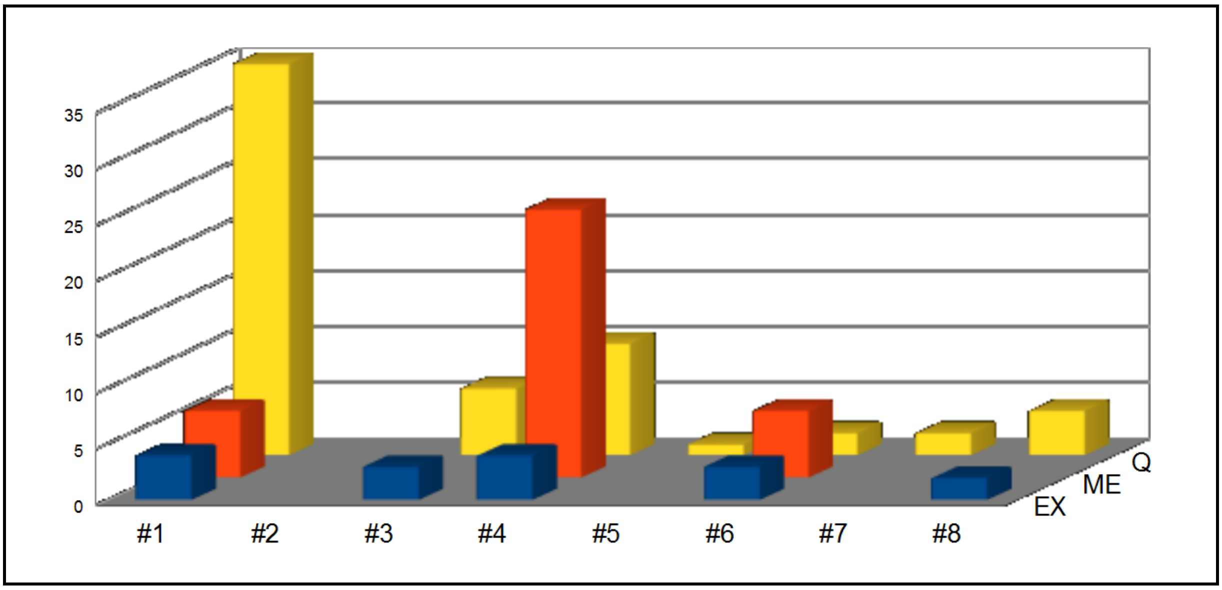

Figure 13*:* Occurrence of execution (EX), misestimation (ME) and quality (Q) errors in the eight sub-segments.

4.4.3 <u>Discussion</u>

4.4.3.1 Error classification

The error classification system by Hughes et al. (2013) developed for more severe cases of stroke impairments could only partially be used in this study, since the patients tested in the TUMLab were mainly chronic stroke survivors and therefore had already received some sort of specific rehabilitation. So only 9 of the 12 error-classes were observed and there were no fatal errors executed by the subjects. Even controls committed some errors. This might be due to their advanced age and the test situation, which stroke survivors are commonly better used to.

RBD patients showed in average 2.94 times as many errors as controls per trial and LBD patients 1.73 (Fig. 4 & 5). The larger number of errors in RBD patients were usually linked to sequence production and apparently, the task of tea-making is quite demanding for patients with ADS symptoms. Interestingly LBD patients did not show a special pattern of errors as one might expect from the fact that LBD patients frequently suffer from apraxia. The reason might be that their symptoms were mild compared to the RBD patients or that the tea-making task is not that challenging for patients with symptoms of apraxia.

4.4.3.2 Error recognition

Among the three error types analyzed, only the misestimation error showed a slight, kinematic peculiarity of the mean peak velocity when segment 4 (‘pour the heated water into the mug’) was affected by an error. In the 4 out of a total of 24 trials the mean velocity was increased but neither path length increased nor movement time decreased. So error trials of this sub-segment must have been performed in a more discrete way. Since this feature appeared in all groups, the source of the observed behavior may not be necessarily related to brain damage but more probably to the advanced age of the subjects. A possible explanation could be a higher weight of a filled kettle in error trials that forces some subjects to move the container in a more stepwise way.

4.4.3.3 Error prediction

For the three error types no clear kinematic peculiarity was observed in the sub-segment preceding an error-affected sub-segment. This might change with upcoming analyses of the other 6 error types, since sequence errors might have preceding phases of confusion combined with longer movement times but this is still to be analyzed. Independently from kinematics, the probabilities of a specific error to occur in a particular segment can be taken into account for error prediction, since misestimation and quality errors occur more frequently in particular sub-segments.

4.4.3.4 Conclusion

We did not find strong indications of successful error recognition and prediction based on monitoring the subjects hand kinematics in these first results. Kinematics may however be able to support error recognition in specific cases. Nevertheless analyses of the 6 remaining error types classes and may provide useful results for errors linked to sequencing.

4.5 Relationship between action errors in AADS and gaze

4.5.1 Results

4.5.1.1 Performance of healthy control subjects

Figure 14 shows the fixation times on the objects of the task of tea-making in healthy controls. Overall fixation times during the different sub-segments show characteristic and expectable gaze behavior (Fig.6). For the analysis of gaze, the time windows for analyzing each segment was shifted forward in time since the eyes typically lead the hand. A gaze-kinematics shift of 0.61 seconds in the beginning and 0.56 in the end respectively has been used according to Land & Hayhoe (2001).

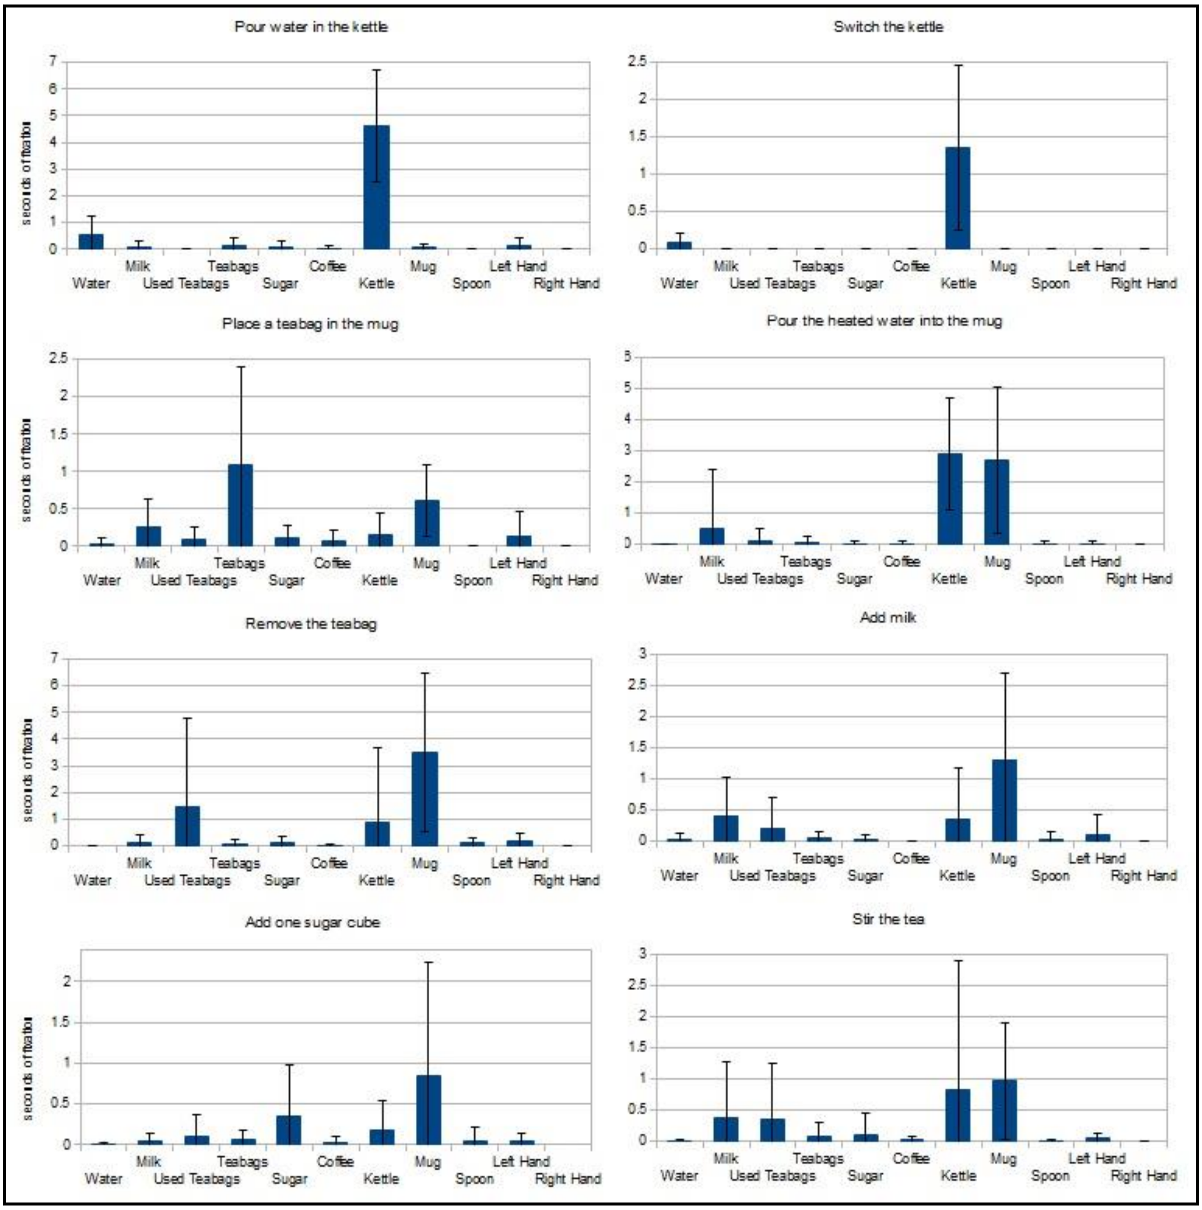

Figure 14: Fixation times in the different sub-segments during the task of tea-making in healthy controls.

In segment 1 ('pour water in the kettle') the control group was mainly fixating the task relevant objects, the kettle and the water container. Other objects were also fixated, probably for locating purposes since the task is in 98% of the trials started with this segment.

In segment 2 ('switch the kettle on') the longest fixations were on the kettle and only a fraction of a second on the water container. This can be considered as a look-back strategy, as the preceding segment's focus was the water container, which is needed to fill the kettle.

In segment 3 ('place a teabag in the mug') fixations were mainly on the teabags but with a high variability. Gaze on the mug, which is the goal object of this segment, was consistent observed in that segment. The long fixations on the teabags can be explained by the entangling of the teabags, which was difficult, especially in the unimanual condition of the task. Segment 3 also contained short fixations on other objects such as the coffee, the milk or the sugar. This behavior may be associated with relocating intentions for planning. Gaze on the left hand is noticeable and might be due to insecurity of acting with the non-dominant hand.

In segment 4 ('pour the heated water into the mug') mainly the kettle and the mug were fixated. A short period of fixating the milk at this point of the task is noticeable. Since segment 5 ('remove the teabag') is sometimes left out, segment 6 ('add milk') is in 12% of the trials the sub-segment following segment 4 and fixations of the milk may therefore reflect further action planning.

Segment 5 ('remove the teabag') shows suitable times of object fixations with the prominent fixation time on the mug. This is most probably based on the fact that the teabag is taken from the mug which contains hot water.

In segment 6 ('add milk'), subjects fixated various objects with the longest fixation on the mug and a shorter on the milk. Note the fixations on the left hand and on the container for the used teabags. Since the milk was located on the left of the subject, it was commonly grasped with the left hand in the bimanual and only in the unimanual right hand condition with the right hand. This guiding gaze was presumably based on a weaker motor control of the non-dominant hand in the mostly right handed subjects. Fixating the container for the used teabags may be considered as look-back, since segment 6 follows segment 5 in 56% of the trials. This quite prominent fixations might be due to assuring that the task is almost complete, as in 49% of the trials segment 8 ('stir the tea') follows segment 6 as a finishing segment and in 12% of the trials they directly finish with segment 6.

In segment 7 ('add one sugar cube') the fixation times describe a typical grasping and dropping action with rather short fixations on the sugar cubes and longer fixations on the mug the latter being the goal of the action.

Segment 8 ('stir the tea') shows a wide distribution of fixated objects. Although the mug receives the highest attention almost all other allocated objects are also fixated. Given that this is the last task segment when preparing the tea, controls fixate already used objects to reassure that they performed well and did not leave out any important steps.

For the first segments, one reason for continuous fixations of the kettle in almost all sub-segments may be that when waiting for the boiling of the water, subjects usually start to perform other sub-segments (e.g. placing a teabag or sugar into the mug), but are repeatedly checking the kettle for the temperature on the display or via the steam coming out of the kettle.

4.5.1.2 Performance of CVA patients

Fixation Times

Figure 15 shows the fixations times of the control group and of the patients in the sub-segments of the tea-making task. The patients produced longer fixation times in all sub-segments but segment 8 ('stir the tea') (Figure 14). This might be due to less reassuring whether they performed correctly.

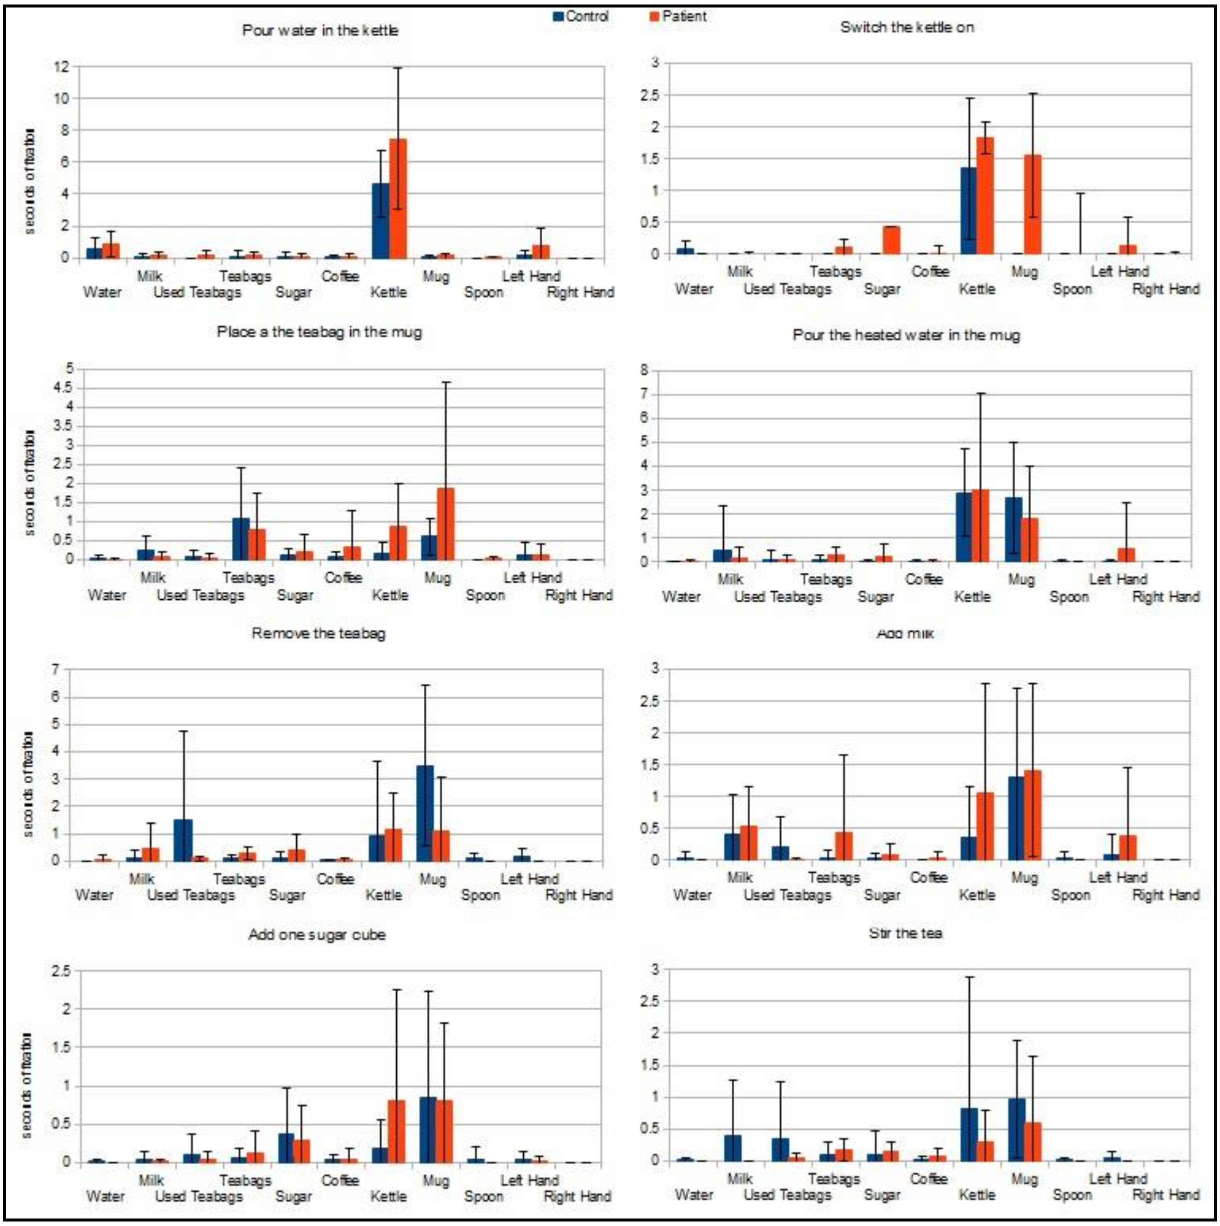

Figure 15: Seconds of fixation during the performance of the tea making segments in patients compared to healthy controls

While both groups showed a comparable gaze pattern for the water container in segment 1 ('pour water in the kettle'), the fixation time of the kettle in patients with approximately 8 seconds is clearly increased. It seems that patients needed more time for checking and guiding the action of filling the kettle.

While in segment 2 (‘switch the kettle on’) the controls only focused on the kettle and the water container but not on the sugar, the teabags or the mug, the patients fixated these objects, especially the mug. The next segments are usually segment 3, 4 and 7, which are using the teabags (3), the kettle (4) and the sugar (7) and all these segments end with putting something into or filling the mug. These looks can therefore be considered as look-aheads for ongoing action planning that seems aggravated or impaired when contrasting to the control group.

In segment 3 (‘place a teabag in the mug’) patients fixated the mug longer than controls. This could indicate that dealing with the swinging of the teabag when placing it into the mug was particularly difficult for the patients. Patients also fixated the sugar and the coffee that are both in containers of the same kind as the teabags. This could be a hint for an impaired locating mechanism in the beginning of the task and/or during object recognition.

In segment 4 (‘pour the heated water into the mug’) both groups fixated the kettle longer than the mug, which seems surprising for a filling action. This could have been due to hesitations because the kettle was quite heavy for an elderly person and contained hot, steaming water. Also, patients looked at their left hand, possibly indicating insecurity and efforts to control the impaired hand (in RBD patients).

In segment 5 (‘remove the teabag’) there was almost no fixation on the used teabags in the patient group in comparison with the control group. It is also worth to be mentioned that patients exhibited longer fixations of the sugar and the milk than of the container for the used teabags in this segment. Since CVA patients finished the task with segment 5 in 44% of the trials and continue with segment 6 in 33% and with segment 7 in 11% of the trials these fixations can be considered either look-backs for assurance of a successful task completion or look-aheads for the upcoming segments 6 or 7.

In segment 6 (‘add milk’) fixations of the teabags and the left hand can be observed in the patient group but not in the control group. While the milk not surprisingly catches most of the attention in both groups, the patient group once more shows fixations of the left hand. Problems in motor control were probably responsible for the prolonged fixations of the left hand in this group. Gaze on the teabags is most probably a look-back as in 33% of the trials segment 5 precedes 6.

In segment 7 (‘add one sugar cube’) the patient group shows longer fixations of the kettle. This can be explained by a different order of sub-segments in the patient group. The controls usually add the sugar after filling the mug with water while patients in several trials add the sugar beforehand.

Segment 8 (‘stir the tea’) has fixations on the coffee, the sugar and the teabags. This again is a hint for an impaired locating mechanism in the beginning of the task and / or during object recognition. Note the shorter fixation times on the kettle, the milk and the used teabags. All these objects are allocated on the peripheral working surface for the tea-making task.

In conclusion, the fixation patterns of the patient group show more inappropriate fixations and differ from the controls. Also, patients show more fixations of their left hand, presumably due to problems in motor control.

4.5.1.3 Gaze and error

For the analysis of errors in the context of gaze, so far 3 classes of errors have been used:

- Execution errors (EX)
- Misestimation errors (ME)
- Quality errors (Q)

For the exploration of the relationship between performance errors and gaze behavior, the fixations times on the objects in the affected and the preceding sub-segment were analyzed qualitatively.

Execution errors

In two of three segments with execution errors, patients showed unnecessary fixations of the teabags, both subjects being RBD patients. The execution errors happened in segments with pouring actions (4 & 6). This is an interesting finding but needs more than two cases to be indicative of a regular association. Overall only 16 execution errors were observed in a total of 67 trials.

Misestimation errors

Analyzing the 36 misestimation errors in the context of gaze, 5 trials showed peculiarities in the affected segment. 4 trials showed peculiarities in the preceding segment. In the affected segments patients had either missing fixations of the container to be filled (mug or kettle), prolonged fixation times on the kettle (in one case almost 12 seconds) or unnecessary fixations of the teabags (slightly shorter than a second). In the preceding sub-segments, gaze was on the left hand for almost 7 seconds in one case and in 4 trials on momentarily irrelevant objects, mainly on the coffee (distractor item), the sugar and the teabags. In 4 out of these 5 trials the affected sub-segment was segment 4 ('pour the heated water into the mug') which is usually followed by segments with a more variable order like adding milk or sugar, removing the teabag or stirring the tea.

Quality errors

In total there were 64 quality errors in the analyzed 67 trials of patients and controls. Of these, 11 showed peculiarities in the gaze behavior, most of them in RBD patients. In some cases patients showed fixations of their left hand or of all ingredient objects, like the coffee, the sugar or the milk, and particularly the teabags and the coffee, also in the preceding sub-segment. Quality errors of control subjects were on the other hand associated with a low number of fixations of task relevant objects like the water container when filling the kettle.

Discussion

Execution errors in the patients probably result from deficits of motor control or from a reduced attentional focus on the current action. This reduced focus can be recognized by fixations on momentarily irrelevant objects.

Errors of misestimation can take place if the filling level of a container is not checked. This was the case in 4 out of 5 trials with gaze peculiarities. But even if the filling level is checked, a reduced focus on the current task can lead to misestimations. This is reflected by unnecessary fixations of the teabags or fixation times reaching up to 12 seconds, resulting from high demands for handling of the kettle.

Quality errors can occur due to an incorrect or careless handling of an object, deficient motor performance due to an impaired hand or an incomplete object preparation for a following action (e.g. the teabag is still in the mug when the subjects starts stirring the tea and therefore the spoon gets entangled). Symptoms of these causalities are observable in the fixation of the left hand, missing fixations of a task relevant object and fixations of the distractor item (coffee).

4.5.2 Conclusion

Gaze patterns are characteristic for the different sub-tasks in the segments. These patterns are however not constant for a certain segment, but can vary substantially within and between subjects. In addition, the uniqueness of a pattern depends on the different sub-tasks. In addition, stroke patients and age-matched control subjects exhibit partly different gaze patterns although the difference is not necessarily related to errors. For example, patients tend to look frequently to the impaired hand. Therefore, action recognition based on gaze information would not be 100% precise. Nevertheless informing models like the Hidden Markov Models employed in the CogWatch system about gaze would almost certainly increase the precision of automatized action recognition, in particular if the models are tolerant to variability. Another advantage of gaze information is that the pattern would be up to a certain degree independent of the spatial position of the objects that may be moved around in natural conditions of daily activities.

It also seems that fixation patterns and fixation times have the potential to support the recognition and prediction of errors in the performance of the tea-making task. Since gaze leads action, fixations could be used to predict following actions and even errors. Implementing eye-tracking and an automatized fixation analysis could therefore provide a valuable contribution in recognizing and predicting errors.

Employing the method in the CogWatch system is not possible at the moment since reliable online gaze recognition is not available yet. There are however strong technical developments in the fields of eye-movement recordings and computer vision that may make this technology available soon.

5. CONCLUSIONS

This report has been concerned with predictive models in the CogWatch project. The first part of the report, sections 3.3 described the technology that underpins the development of the task model (TM) in P1.2, for the tea-making task, and which will be included in P2 for the tooth-brushing task. The second part, section 4, described psychological experiments aimed at measuring the extent to which knowledge of body kinematics and gaze can be used to predict errors in the tea-making task.

In CogWatch, the purpose of the TM is to monitor a user's progress through the task, to detect errors and to provide sufficient information to the CogWatch system for useful cues to be created. The rationale for choosing a Markov Decision Process (MDP) based approach to task modelling was presented in deliverable D3.3.1. An MDP-based TM for tea-making has been developed at UOB using a hierarchical description of the task. The TM is implemented in Python and has been fully integrated into the C# environment of the CogWatch system. In addition, a program was written to simulate a user of the system (the 'SimU') based on statistics of sub-goal sequences observed in trials of patients and healthy controls. The SimU allowed a substantial number of 'experiments' to be conducted to evaluate the TM. The evaluations are in terms of the user task completion rate as a function of the accuracy of the action recognition (AR) system, and the compliance of the user (i.e. whether the user follows the cues created in response to the TM outputs). For example, with an AR error rate of 10% the TM achieves a user completion rate greater than 90%.

Some degree of AR error is inevitable. Therefore the more recent research focussed on developing a TM that is robust against such errors. The new TM uses a Partially Observable (MDP). The key difference between the MDP and POMDP TMs is that while the MDP maintains a single estimate of the state of the user in the task, the POMDP 'belief state' is a distribution over all of the MPD states. Experiments have shown that the POMDP-based TM can support user task completion rates of 90% with AR error rates greater than 20%.

Section 4 describes experiments conducted at TUM to measure the relationship between action errors in the tea-making task and user kinematics and gaze. The objective is to provide the psychological basis for incorporating this type of information for error prediction in future TMs. Sixty-seven trials were conducted, of which 41 were performed by controls, 16 by patients with right brain damage (RBD) and 10 by patients with left brain damage (LBD). During trials subjects wore a SMI-ETG eye tracking device and their movements were tracked using the Qualysis system. The data streams were synchronized, segmented into sub-goals, and errors were categorized according to the scheme proposed by Hughes et al. (2013). Analyses were conducted on segments exhibiting errors for error detection, and on the previous segment for error prediction.

The main conclusions of the study are that error recognition and prediction based on monitoring the subjects hand kinematics does not appear to be viable. However, fixation patterns and fixation times do have the potential to support the recognition and prediction of errors. Since gaze leads action, fixations could be used to predict following actions and even errors. Incorporation of eye-tracking and an automatized fixation analysis into the TM could provide a valuable contribution in recognizing and predicting errors.

REFERENCES

De Kleine, E., Van der Lubbe, R.H., “Decreased load on general motor preparation and visual-working memory while preparingfamiliar as compared to unfamiliar movement sequences.”, Brain Cogn. 2011, 75(2):126-34, 2011.

Graybiel, A.M., “The basal ganglia and chunking of action repertoires.” Neurobiol Learn Mem. Jul-Sep;70(1-2):119-36, 1998.

Hermsdörfer, J., Bienkiewicz, M., Cogollor, J.M., Russell, M., Jean-Baptiste, E., Parekh, M., Wing, A.M., Ferre, M., Hugues, C., “CogWatch: Automated Assistance and Rehabilitation of Stroke-induced Action Disorders in the Home Environment”, HCI Aspects of Optimal Healing Environments, 2013.

Hughes, C., Baber, M., Bienkiewicz, M. & Hermsdörfer, J. “Application of Human Error Identification (HEI) Techniques to Cognitive Rehabilitation in Stroke Patients with Limb Apraxia”. HCI, 8 (2013), 463-471, 2013.

Humphreys, G. W. & Forde, E. M. E. “Disordered action schema and action disorganisation symdrome”. In Forde, E. M. E., Rusted, J., Mennie, N., Land, M. & Humphreys, G. W. (2010). The eyes have it: An exploration of eye movements in action disorganisation syndrome. Neuropsychologia 48, 1895-1900, 1998.

Jean-Baptiste, E.M.D, Nabiei, R., Parekh, M., Fringi, M., Drozdowska, B., Baber, C., Jancovic, P., Rotshein, P. and Russell, M. “ Intelligent Assistive System Using Real-Time Action Recognition for Stroke Survivors”. In Proceedings of the IEEE International Conference on Healthcare Informatics (ICHI 2014).

Land, M. F. & Hayhoe, M., ”In what ways do eye movements contribute to everyday activities?”, Vision Research 41, 3559-3565, 2001.

Levin E., Pieraccini R., Eckert W., “A stochastic model of human machine interaction for learning dialogue strategies.” IEEE Transactions on Speech and Audio Processing: Vol.8, No. 1., January, 11-23, 2000.

Pflügler, J., Schlegel, A., Jean-Baptiste, E.M.D., Rotshtein, P., Pastorino, M., Rojo, J., Cogollor, J.M., Arredondo, M.T., Bienkiewicz, M. and Hermsdörfer, J. **“**Using human-computer interface for rehabilitation of activities of daily living (ADL) in stroke patients. Lessons from the first prototype**.”** In Proceedings of the International Conference on NeuroRehabilitation (ICNR)2014

Williams, J.D., Poupart, P., Young, S., “Factored partially observable Markov Decision Processes for dialogue management”. Computer Speech and Language: p. 393-422, 2005.

Young, S.J., Gasic, M., Keizer, S., Mairesse, F., Schatzmann, J., Thomson, B., Yu, K: “The Hidden Information State model: A practical framework for POMDP-based spoken dialogue management**”,** Computer Speech & Language 24(2): 150-174 (2010)

Williams, J.D., Poupart, P., Young, S.J., “Partially Observable Markov Decision Processes with Continuous Observations for Dialogue Management”. Recent Trends in Discourse and Dialogue. Volume 39, 2008, pp 191-217