

# A Dynamic Decision Network Framework for Online Media Adaptation in Stroke Rehabilitation

YINPENG CHEN, WEIWEI XU, HARI SUNDARAM, THANASSIS RIKAKIS, and SHENG-MIN LIU  
Arizona State University

In this article, we present a media adaptation framework for an immersive biofeedback system for stroke patient rehabilitation. In our biofeedback system, media adaptation refers to changes in audio/visual feedback as well as changes in physical environment. Effective media adaptation frameworks help patients recover generative plans for arm movement with potential for significantly shortened therapeutic time. The media adaptation problem has significant challenges—(a) high dimensionality of adaptation parameter space; (b) variability in the patient performance across and within sessions; (c) the actual rehabilitation plan is typically a non-first-order Markov process, making the learning task hard.

Our key insight is to understand media adaptation as a real-time feedback control problem. We use a mixture-of-experts based Dynamic Decision Network (DDN) for online media adaptation. We train DDN mixtures per patient, per session. The mixture models address two basic questions—(a) given a specific adaptation suggested by the domain experts, predict the patient performance, and (b) given the expected performance, determine the optimal adaptation decision. The questions are answered through an optimality criterion based search on DDN models trained in previous sessions. We have also developed new validation metrics and have very good results for both questions on actual stroke rehabilitation data.

Categories and Subject Descriptors: J.3 [Life and Medical Sciences] Health; H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces—*Computer-supported cooperative work*; I.6.4 [Simulation and Modeling] Model Validation and Analysis

General Terms: Algorithms, Experimentation, Human Factors

Additional Key Words and Phrases: Biofeedback, media adaptation, dynamic decision network, mixture of experts

## ACM Reference Format:

Chen, Y., Xu, W., Sundaram, H., Rikakis, T., and Liu, S.-M. 2008. A dynamic decision network framework for online media adaptation in stroke rehabilitation. *ACM Trans. Multimedia Comput. Commun. Appl.* 5, 1, Article 4 (October 2008), 38 pages. DOI = 10.1145/1404880.1404884 <http://doi.acm.org/10.1145/1404880.1404884>

## 1. INTRODUCTION

The goal of this article is to develop a Dynamic Decision Network (DDN)-based media adaptation framework for use in a multimodal biofeedback system [Chen et al. 2006] for stroke patient rehabilitation. Stroke rehabilitation is an important problem—every 45 sec. someone in the United States suffers a stroke, often leading to physiological impairment. Effective media adaptation frameworks can potentially lead to significantly shortened therapeutic procedures (typically taking years). In this paper, the media adaptation refers to changes in audio/visual feedback as well as changes in physical

Author's address: H. Sundaram; email: Hari.Sundaram@asu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2008 ACM 1551-6857/2008/10-ART4 \$5.00 DOI 10.1145/1404880.1404884 <http://doi.acm.org/10.1145/1404880.1404884>

ACM Transactions on Multimedia Computing, Communications and Applications, Vol. 5, No. 1, Article 4, Publication date: October 2008.

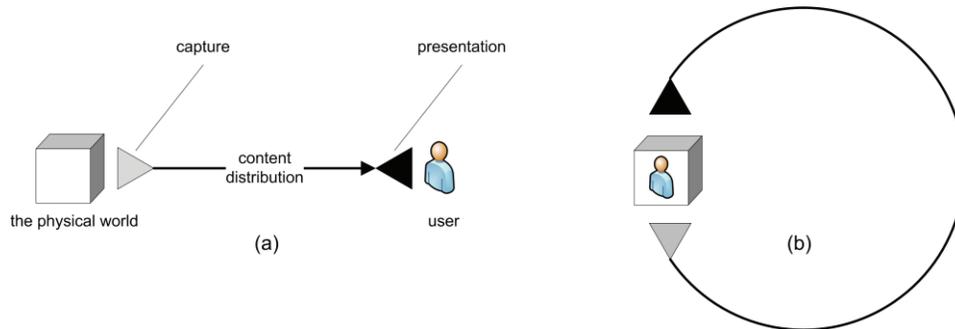


Fig. 1. Contrast between conventional multimedia systems (a) and experiential media systems (b). In the former, the capture, production, distribution and consumption of media is asynchronous, and non colocated (e.g., a film). The producer of content has a significant role in the interpretation of the media. In the latter, the production and consumption of media is synchronous, colocated and the semantics are a joint construction between the user and the system.

environment. The mediated rehabilitation system is an example of an experiential media system [Sundaram and Rikakis 2006].

### 1.1 Context: Experiential Media Systems

Our interaction with our physical surroundings provides us with a complex multimodal experience. For example, let us examine the simple task of reading a book. In this simple, everyday experience, the reader experiences the weight of the book, the texture of the paper, and the temperature of the book; executes coordinated hand-eye movements to turn the page; and decodes the visual symbols imprinted on a page. The task of reading a book involves sensing, perception, and cognition, as well as action that alter the physical world (turning the page). It is a task that requires manipulation of data in a control loop, and this can occur at multiple time scales. In addition to requiring real-time sensing, perception and decision making, the interaction with the physical world serves as the primary mechanism for us to acquire new knowledge.

Experiential media systems [Sundaram and Rikakis 2006], refer to real-time, physically grounded multimedia systems that enable human beings to acquire knowledge through interaction with a physical system. An important measure of success of such systems is the ability of the participants to transfer (or generalize) the learned knowledge to physical interactions outside of the system, in the physical world. These systems involve participating in our multi-sensory environment through the use of real-time, context-aware computational elements—sensors and presentation devices that are part of the physical environment.

Experiential media systems involve a bi-directional transfer of semantics. Multimedia content analysis primarily focuses on difficult problems in the offline analysis and labeling of captured data, for example, labeling of a digital photo as “outdoor”/“Jane,” summarizing produced media from television or films etc. [Sundaram and Chang 2000; Sun et al. 2002; He et al. 2003; Ma and Zhang 2003]. Here, the semantics are inferred from the media artifacts that document human activity (e.g., a photograph of wedding). However, in an experiential media system, semantics are inferred not just by the system which observes human activity, and which begins to parse meaningful patterns. The human participant is immersed in a real-time audio-visual environment, wherein the system has embedded semantics. Hence the semantics that emerge through such interaction are a *joint* construction by the user and the system. Figure 1 shows the contrast between conventional multimedia systems and experiential media systems.

The successful development of such systems requires us to integrate knowledge from multiple disciplines—computational elements from science and engineering, ability to develop meaningful immersions from the arts and understanding cognition and learning from psychology and education. The development of the computational elements of experiential media systems is challenging and requires knowledge from several disciplines within engineering—artificial intelligence (robotics [Brooks 1991a, 1991b; Brooks et al. 1997]), human-computer interaction (tangible interfaces [Ishii and Ullmer 1997; Ishii et al. 1998; Ullmer and Ishii 2000; Mazalek et al. 2002], embodied interaction [Dourish 2001]), distributed cognition [Hutchins 1995; Kirsh 1995; Hollan et al. 2000] and ubiquitous computing [Weiser 1993; Abowd et al. 2002]. Each of these fields has led us to reimagine computation integrated with the physical world. Each focuses on a complementary aspect of the interaction with the physical world, however this knowledge is not well integrated with research in multimedia computing.

In this paper, we shall discuss our efforts to develop adaptation frameworks for a specific type of an experiential media system—a stroke patient rehabilitation system [Chen et al. 2007]. In this system, the knowledge that is to be transferred is restricted to a simple, but important functional task—to reach out and grasp a cup. This simple task is enormously challenging for stroke patients, who have lost the ability to control their arm, and hence must be retrained (i.e., they need to relearn the necessary motor control) to perform the task. Since experiential media systems are integrated with the physical world, the parameter space that is required to adapt the system, so as promote optimal transfer of knowledge, is very large (~250 parameters in our system). Determining which parameters need to be changed, online, is therefore a highly challenging task.

## 1.2 Related Work

There has been prior work on biofeedback therapy. Virtual Reality (VR) is an emerging and promising technology for task-oriented biofeedback therapy to improve motor function in individuals with stroke and other neurological disorders [Holden and Dyar 2002; Gallichio and Kluding 2004]. VR can provide an effective human-computer interface that allows users to interact with a complex, highly detailed and engaging multimodal feedback in response to physical actions. This has significant potential in augmenting the traditional task-oriented therapy training. It has been shown that task learning in a VR can be transferred into real world task performance [Jack et al. 2001]. Holden et al. [1999] utilized VR to train reaching and hand orientation of stroke patients by developing a virtual mailbox environment with different slot heights and orientations. The design and implementation of a virtual kitchen used to practice common daily-living activities is proposed in White et al. [2005]. It is shown in [Holden and Dyar 2002] that nine recruited participants achieve significant improvement in the Fugl-Meyer (FM) score, the Wolf Motor Function (WMF) score, and selected strength tests after VR-based training. Using Cyberglove or the RMII glove [Merians et al. 2002], investigators showed that the patients improved grasping force, finger joint range of motion, and movement speed after two weeks of VR-based biofeedback therapy.

There has been extensive prior work on using dynamic models in multimedia [Xie et al. 2003] as well as in computer vision [Murphy et al. 2003]. The Hidden Markov Model (HMM) has been proved quite adapted in domains such as speech recognition [Rabiner 1989], modeling human activity [Liao et al. 2004], probabilistic plan recognition [Bui 2003] and robot action [Fox et al. 2006]. The inference and learning in Dynamic Decision Network (DDN) have been investigated in Kjaerulff [1995], Darwiche [2001], and Murphy [2002].

Our research bears resemblance to a real-time feedback control problem. There, the most relevant work appears in robotic control/planning problems. Extensive summaries of the robot planning literature are available in [Latombe 1991; LaValle 2006]. In [Burridge et al. 1999], the idea of dynamic tasks by sequential composition of feedback primitives was presented. In addition, Grupen and Coelho [2002]

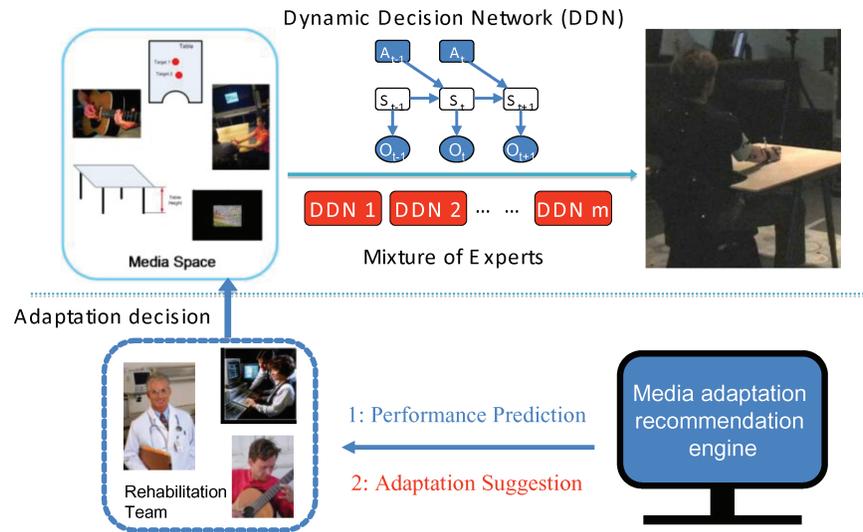


Fig. 2. Human-computer joint media adaptation framework. The adaptation decisions are made by rehabilitation team that includes therapists, doctors, artists and engineers. Our media adaptation recommendation engine provides the rehabilitation team two kinds of useful information: performance prediction and adaptation suggestion. In our adaptation recommendation engine, we model the relationship between the media adaptation and the subject movement by using mixture-of-experts-based Dynamic Decision Network (DDN).

proposed a system that learns optimal control policies in an information space that is derived from the changes in the observable modes of interaction between the robot and the object it is manipulating. There is a substantial body of previous work applying Partially Observable Decision Markov Processes (POMDP) [Simmons and Koenig 1995; Kaelbling et al. 1998; Pineau et al. 2003; Smith and Simmons 2004; Theodorou and Kaelbling 2004; Theodorou et al. 2004; Spaan and Vlassis 2005] to mobile robot navigation. At the beginning, Simmons et al. [Simmons and Koenig 1995] proposed a simple heuristic solution. Since it is computationally intractable to compute the exact optimal finite or infinite-horizon solution of a POMDP, sophisticated approximation algorithms [Pineau et al. 2003; Theodorou and Kaelbling 2004; Spaan and Vlassis 2005] were proposed. These algorithms take advantage of the fact that the set of states that are reachable under a reasonable control policy is dramatically smaller than the original space. Theodorou et al. [2004] used hierarchical POMDPs to represent multi-resolution spatial maps for indoor robot navigation. Hsiao et al. [2007] addressed the finger grasping problem by using POMDP for robot manipulation in.

### 1.3 Our Approach

We present a human-computer joint adaptation framework (see Figure 2) in this paper which is built upon our recent work [Chen et al. 2007]. In this human-computer joint adaptation framework, the adaptation decisions are made by rehabilitation team that includes therapists, doctors, artists and engineers. Our proposed media adaptation recommendation engine helps the rehabilitation team for adaptation decision by providing two kinds of useful information: (a) performance prediction and (b) adaptation suggestion.

The central insight in our adaptation solution is the realization that there exists a *duality* between the robot navigation problem and the stroke patient rehabilitation problem. In the former, the robot is completely controllable, and must operate within an uncontrollable natural world. The goal for the

robot is to learn an optimal plan, to navigate the physical world. In our problem the environment in which the patient is immersed is completely controllable; however, the behavior of the stroke patient is not within our control. The goal *for the system* is to transfer motor-control semantics related to the functional task (reach and grasp). This duality causes solutions that span seemingly very different problems, share the same mathematical structures. Our solution is similar to the POMDP.

The media adaptation problem has four significant challenges: (a) the joint subject movement and media adaptation is high dimensional, making model learning difficult, (b) the patient performance across sessions is highly variable due to external factors such as medication, lack of proper sleep, (c) the performance within session is affected by physiological factors such as being tired and (d) the physical therapists/doctors and artists plan two to three steps ahead—this is a non first-order Markov process, making the learning task hard.

The main contribution of this article is in the development of a mixture-of-experts-based Dynamic Decision Network (DDN) to model the relationship between the media adaptation and the subject movement. The subject functional task is to reach for a target. The decision network contains three types of nodes—adaptation decision ( $A_t$ , observable), state describing the movement plan ( $S_t$ , hidden) and subject movement observations ( $O_t$ , observable). The decision node  $A_t$  represents the media adaptation decision after subject movement observations at time  $t$ . We make decision on *changes* rather than on values to make the DDN parameter learning problem tractable. The state node  $S_t$  represents the subject hidden state in reaching/grasping plan acquisition. The semantics of the hidden state represent the “goodness” of the body plan (in movement sequence) to reach the target. The observation node  $O_t$  represents the subject’s reaching/grasping performance in the set  $t$ . We train DDN mixtures per patient, per session. In online adaptation, we dynamically search for a proper DDN model from all DDN models trained in previous sessions and use the selected model to recommend adaptations for the rehabilitation team.

Our framework addresses two issues—(a) performance prediction: given a specific adaptation suggested by the domain experts, predict the patient performance and (b) adaptation suggestion: given the expected performance, determine the optimal adaptation decision. We have very good results for both questions on actual stroke rehabilitation data.

#### 1.4 Organization of the Article

The rest of the article is organized as follows. In the next section we overview the biofeedback system. In Section 3, we discuss the media adaptation problem in the biofeedback system and present the key challenges. In Section 4, we present the adaptation model using the DDN. We propose the performance prediction algorithm and the adaptation recommendation algorithm in the Section 5. In Section 6, we show how to use DDN based adaptation framework in biofeedback system. We show the experimental results in Section 7 and present our conclusions in Section 8.

## 2. BIOFEEDBACK SYSTEM

In this section we summarize our current biofeedback environment for stroke patient rehabilitation [Chen et al. 2006] and discuss limitations for therapist decision making.

There is a significant difference between our first reported work on the biofeedback system [Chen et al. 2006] and the current work. While the earlier paper focused on the *overall design rationale* of the entire biofeedback system, including movement analysis, archiving, feedback generation this paper has a narrow focus on a single aspect of the system. Specifically, we focus on how to best predict the performance of the subject, given certain system parameters, and conversely, suggesting the optimal change of system parameters to help achieve the certain performance targets given by the domain experts.

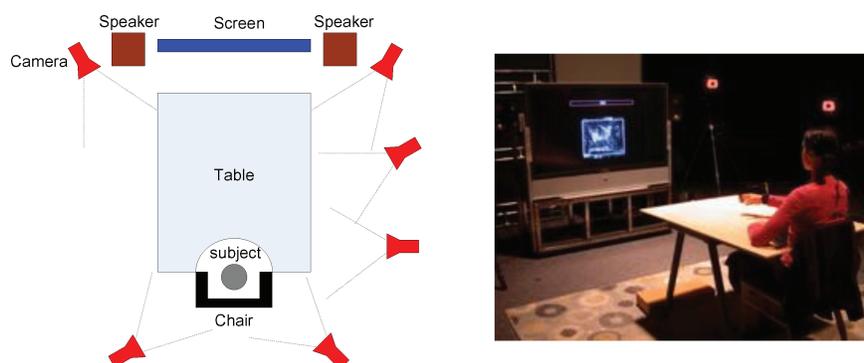


Fig. 3. Physical setup of biofeedback system. Left: physical setup diagram. Right: snapshot of the real environment used by a normal subject.

This section is organized as follows. Next we discuss the physical environment, including how the different components for archiving, analysis and display integrate. Then, in Section 2.3, we show the movement of the subject creates an immersive multimedia feedback. In Section 2.4, we discuss the rehabilitation procedure and conclude with a section on limitations of biofeedback therapeutic process.

## 2.1 Functional Task

The functional task for our patients is *reaching out with the right arm and grasping the target*. This task was chosen because it is the basic arm movement for the stroke patient rehabilitation. This function task includes three subgoals: (a) reach—reach for and grasp the target successfully, (b) open—open the arm to reach for the target without shoulder and torso compensation, and (c) flow—reach for the target smoothly. Therefore, our biofeedback environment encourages subjects to achieve these three subgoals through the engaging and purposeful audio/visual feedback.

## 2.2 Integrating Multimedia Components into the Physical World

We first introduce the physical environment. Figure 3 shows the physical setup of our biofeedback system. The environment need six motion capture cameras, two speakers, one screen, one table and one chair. The subject wears 12 markers on his/her arm, hand and torso and sits at one end of the table. At the other end of the table there are a big screen showing up the visual feedback and two speakers playing audio feedback. Six motion capture cameras are looking over the entire capture volume. Note that the six cameras are asymmetrically placed because we are focusing on the right arm rehabilitation.

We now summarize our current biofeedback system [Chen et al. 2006]. The system integrates seven computational subsystems (see Figure 4): (a) Motion capture; (b) Task control; (c) Motion analysis; (d) Visual feedback; (e) Audio feedback; (f) Media archival and (g) Visualization. All seven subsystems are synchronized with respect to a universal time clock. The motion capture subsystem we are using is produced by Motion Analysis Corporation. We use six near-infrared cameras running at 100 frames per second to track the three-dimensional positions of reflective markers that are placed on the subject. The task control subsystem provides a parameter panel where we can adjust the parameters related to the reaching and grasping task. The real-time motion analysis subsystem smoothes the raw sensing data, and derives an expanded set of task specific quantitative features. It multicasts the analyzed data to the audio, visual and archival subsystems at the same frame rate. The audio and visual subsystems adapt their auditory and visual responses dynamically to selected motion features under different

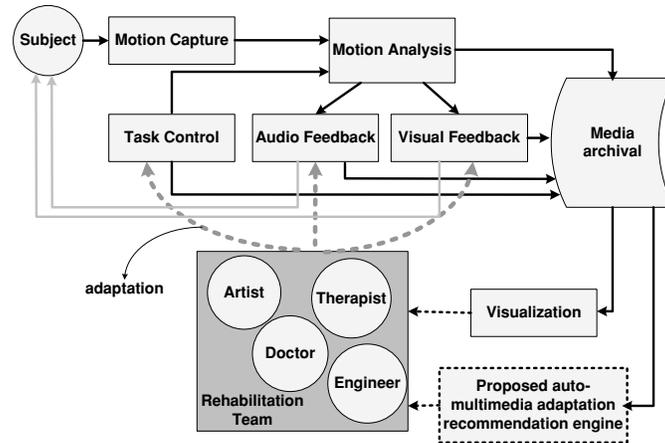


Fig. 4. The flowchart of the biofeedback system.

feedback environments. The archival subsystem continuously stores the motion analysis data as well as the feedback data for the purpose of annotation and offline analysis. The visualization subsystem visualizes the analysis results of the subject's performance [Xu et al. 2006].

The rehabilitation team includes therapists, medical doctors, artists and engineers. The visualization subsystem helps the rehabilitation team to make decisions about how to change the environment for achieving a successful rehabilitation. The system can be fine-tuned through adjusting the parameters in task control, audio feedback engine and visual feedback engine. In this paper, we are proposing an automatic media adaptation recommendation engine (see Figure 4) that provides media adaptation recommendations to be used by the rehabilitation team.

### 2.3 Coupling Movement to Feedback

The structure of the feedback environment and its relationship to the achievement of the goals are based on well established principles regarding the role and function of art [Grout and Palisca 2001]. At the highest level of its structure the environment must communicate to the patient the messages that can encourage the accomplishment of the movement goals. These messages are: *reach, open, flow*.

The overall idea driving the mappings is that spatial and target information is better communicated through visuals and complex time series data are better communicated through audio. The movement parameters allowing successful manipulation of the environment are the key parameters of an everyday reaching and grasping movement. Thus, the environment can be easily connected in terms of action to its goal and does not require specialized movement training. Figure 5 shows the visual feedback under four cases.

The mappings and contents follow a similar structural hierarchy as the movement parameters and goals with submessage levels supporting the communication of each larger message. As is the case of movement parameters, there are feedback parameters that control the feedback generation. The subject can quickly understand the mapping from the arm movement to feedback. Through practice, the subject can control the feedback by moving the right arm to achieve the goals. The detail of movement-feedback coupling is discussed in Chen et al. [2006].



Fig. 5. Visual feedback. Top left: particles begin to form the image as the hand approaches the target. Top right: Image pulled to the right when subject is off target on the right. Bottom left: Vertical bands appear when the subject has wrong target height. Bottom right: completion of image when the subject grasps the target successfully.

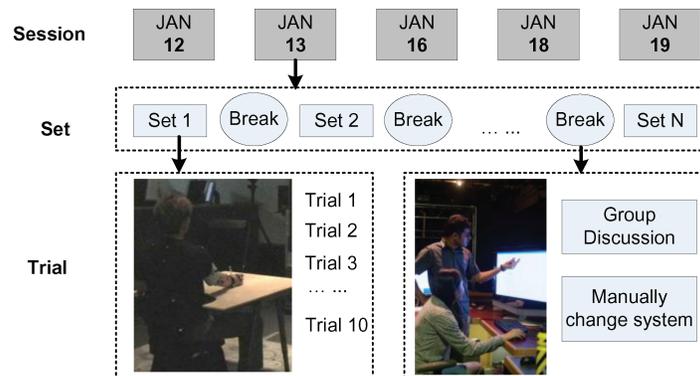


Fig. 6. Rehabilitation procedure. Every session includes several sets. Each set includes ten reaching trials. There is a short break between two consecutive sets. In this break, the rehabilitation team decides how to fine-tune the system after group discussion.

## 2.4 Rehabilitation Procedure

We now introduce the rehabilitation procedure by using our biofeedback system. Let us denote every subject visit as *session*. For each session, there are several sets. In each set, the environmental conditions (physical state, audio and visual parameters) remain fixed. Each set includes ten reaching trials. In each trial, the subject reaches out the right arm toward the virtual target, grasps the target and return back to the rest position. Figure 6 shows the rehabilitation procedure.

The rehabilitation team adapts the system during the short break (typically two minutes) between two consecutive sets. The team discusses the subject’s movement performance through the visualization subsystem which shows the subject’s performance for the previous sets. Then the group decides how to fine-tune the system (e.g., change musical instrument) to help the patient achieving a generative reaching/grasping plan.

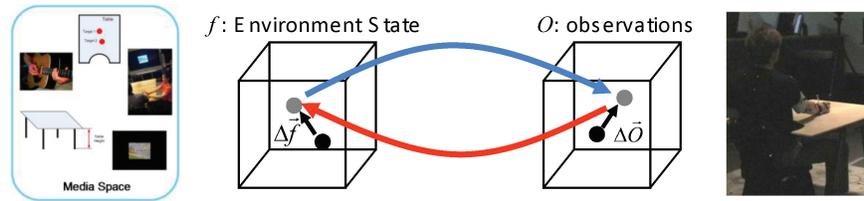


Fig. 7. Media adaptation diagram.

## 2.5 Limitations of Current System

The major limitation of the current system [Chen et al. 2006] is that the media adaptation is totally done by the human experts (rehabilitation team) without accurate data driven analysis. The human experts may miss the information about the relationships between the media adaptation decisions and the patient's movement. This is because the dimensionality of media parameters and movement parameters is very high ( $\sim 100$  parameters) and each rehabilitation takes long time (about 2 hours). Therefore, we are proposing an automatic media adaptation recommendation engine (see Figure 4) that suggests optimal media adaptations to be used by the team. Our proposed adaptation engine is data driven which analyzes the relationships between the subject movement and the media adaptation in previous rehabilitations. This adaptation recommendation subsystem will help the rehabilitation team to make the correct adaptation decision.

## 3. MEDIA ADAPTATION

In this section, we discuss the media adaptation problem in our biofeedback system, and its importance, present the key challenges of the adaptation problem and finally propose our solution and how to apply it in three biofeedback scenarios.

### 3.1 The Problem

The media adaptation problem is:

*How to adapt the biofeedback environment to help the subject acquire a generative plan for reaching and grasping movement?*

Mathematically, the problem is stated as follows: how to find optimal the change in feedback space  $\Delta f$  which results in an expected change in the observation space (patient movement)  $\Delta O$  (see Figure 7). Patient performance is variable both across subjects and within the same subject. The use of media adaptation allows the subjects under different conditions to successfully regain the generative plan for reaching/grasping the target.

The environment adaptations include five parts:

1. reaching/grasping task parameters (e.g., target position)
2. audio feedback parameters (e.g., musical instrument)
3. visual feedback parameters (e.g., the speed of image particles coming together)
4. physical environment parameters (e.g., table height)
5. therapist instruction (e.g., focusing on elbow extension)

We found that in the patient rehabilitation, a combination of these five adaptations works well for the subject to achieve a good reaching/grasping plan. In this paper, we address the automatic data-driven recommendation for adaptation 1–3. We plan to address 4–5 in the future work.

### 3.2 Why Do We Need Adaptation?

The media adaptation is important for several reasons: (a) patients are very different in terms of their capabilities, (b) there is significant performance variability within the same patient across different days (c) we need to change the environment for patients to learn all the variables of the motor plan (d) the environment change helps patient's engagement in rehabilitations. Note that while the adaptation is important, we can *not* change too quickly, as this will affect learnability and patient engagement.

### 3.3 Why Is It Difficult?

There are three key aspects for a successful adaptation: (a) the domain knowledge about the stroke patient, (b) the domain knowledge about the music and visual feedback, and (c) the relationship between the media adaptation and the patient movement. It is difficult to develop a fully automated adaptation framework since the domain knowledge about the stroke patient and arts form (i.e., key aspect (a) and (b)) cannot be easily learned from the data. Hence, we need the rehabilitation team to make adaptation decisions. The rehabilitation team includes doctors, therapists, artists and engineers. However, in the actual rehabilitation, the rehabilitation team also finds that it is very challenging to change the environment because they cannot figure out the relationship between the patient movement and the environment change (i.e., key aspect (c)) due to four reasons:

1. *High dimensionality.* The spaces of both subject arm movement and environment parameters have very high dimensionality. There are more than one hundred parameters related to the multijoint coordinated arm movement training and they are correlated. And there are five groups of environment parameters (see Section 3.1). The combination of environment changes has very high dimensionality and the relationship between the improvement of movement performance and the environment change is complicated and needs to be learnt from the data.
2. *Consistency.* We have had over 30 sessions with the patients and a key observation of stroke rehabilitations is that the subject physiological conditions are different across sessions. This can happen due to many reasons including, taking of medication, not sleeping correctly (i.e., posture) the night before therapy. Hence the rehabilitation team can not make performance continuity assumptions (i.e., patient will move at 10cm/s etc); or make assumption on what specific environment condition is optimal.
3. *Engagement.* It is a challenge for the subjects to remain attentive and motivated during a long and tedious session and they easily become physically and mentally tired. Therefore, the media adaptation framework should also consider how to hold attention through engagement of subjects.
4. *Not 1st-order Markov.* The artists and therapists do not only make changes for the next set, but may make plans for the next several sets (say next three sets). However, learning higher order Markov processes is severely constrained due to limited data. In this paper, we still use the first Markov process to approximate the real-world decision-making process.

Therefore, we design an adaptation recommendation engine to help the rehabilitation team by computing the relationship between the media adaptation and the patient performance. We introduce our solution in the next section.

### 3.4 Our Solution

We present a human-computer joint adaptation framework in our biofeedback system. Our idea is to compute the relationship between the subject movement performance and the environment change and provide useful suggestions for the rehabilitation team. Based on the automated suggestions and the

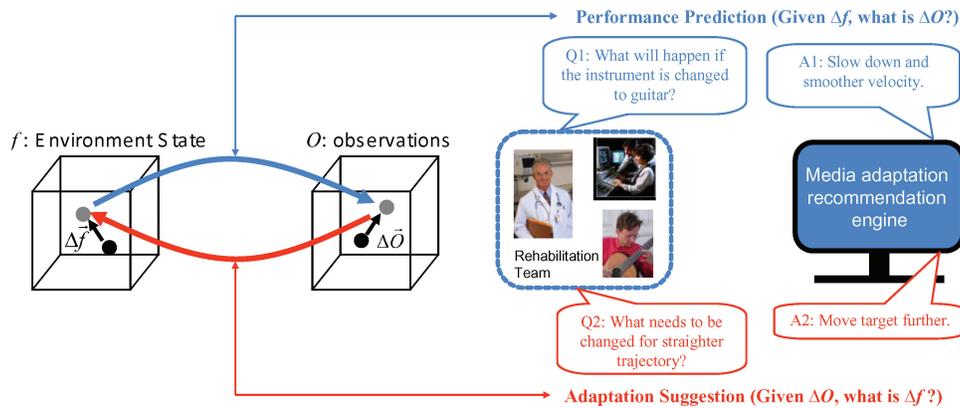


Fig. 8. The diagram of automatic adaptation recommendation for queries of the rehabilitation team.

therapist domain knowledge, the rehabilitation team makes decisions for media adaptation. Figure 8 shows the interaction between the rehabilitation team and the adaptation recommendation engine. The reason why a fully automated adaptation is not attempted is that it is very difficult to account for the extensive experience of the therapist. For example, this experience is very valuable in assessing input at the beginning of the session when the subject state is only available qualitatively e.g., “I do not feel too good today—I slept on my couch last night,” “I took sleep medication last night.” We propose to provide adaptation suggestions by answering the following questions:

- Q1. *Performance prediction.* Given the target environment adaptation  $\Delta f$ , the algorithm predicts the subject movement performance for the next set  $\Delta O$ . (e.g. “what will happen if the instrument is changed to guitar?”)
- Q2. *Adaptation suggestion.* Given the expected subject movement performance  $\Delta O$ , computer adaptation provides the optimal change of the environment  $\Delta f$ . (e.g. “what needs to be changed for straighter hand trajectory?”)

In this paper, computer adaptation only suggests which parameter should change, without considering the quantity of parameter change. The main reason is because of the high dimensionality of the media parameter space and not enough evidence to make accurate predictions. This suggestion is valuable as it dramatically reduces the number of parameters considered by the therapist. The rehabilitation team can then decide the exact quantity.

### 3.5 Three Adaptation Scenarios in Biofeedback System

We apply our adaptation framework (see Section 3.4) on three adaptation scenarios in the biofeedback system: (a) spatial accuracy scenario, (b) hand trajectory scenario and (c) hand velocity scenario. These three scenarios are related to three subgoals (*reach*, *open* and *flow*) of reaching/grasping task (see Section 2.1) respectively. The spatial accuracy scenario is related to the *reach* subgoal. In this scenario, the adaptation goal is to help subjects reach close to the target (see Figure 9 Left). The hand trajectory scenario is related to the *open* subgoal. The adaptation goal is to help subjects straight hand trajectory (see Figure 9 Middle). The hand velocity scenario is related to the *flow* subgoal. The adaptations in this scenario attempt to make the hand velocity smooth (see Figure 9 Right). These three scenarios are selected by therapists and bioengineering experts. We discuss the adaptations in three scenarios in detail later in this paper (see Section 6).

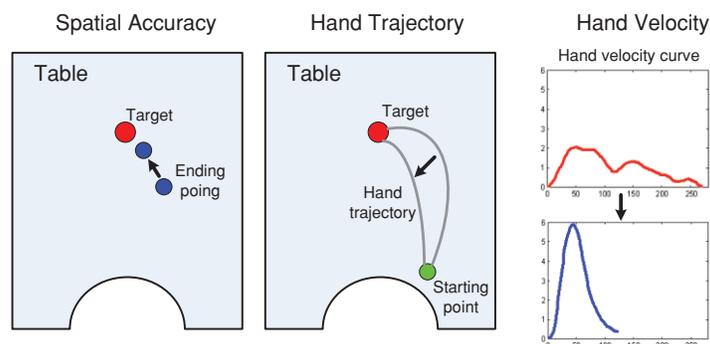


Fig. 9. Three adaptation scenarios in the biofeedback system. Left: spatial accuracy scenario. The goal is to help subjects reach close to the target. Middle: hand trajectory scenario. The goal is to help subjects straight hand trajectory. Right: hand velocity scenario. The adaptations attempt to make the hand velocity smooth.

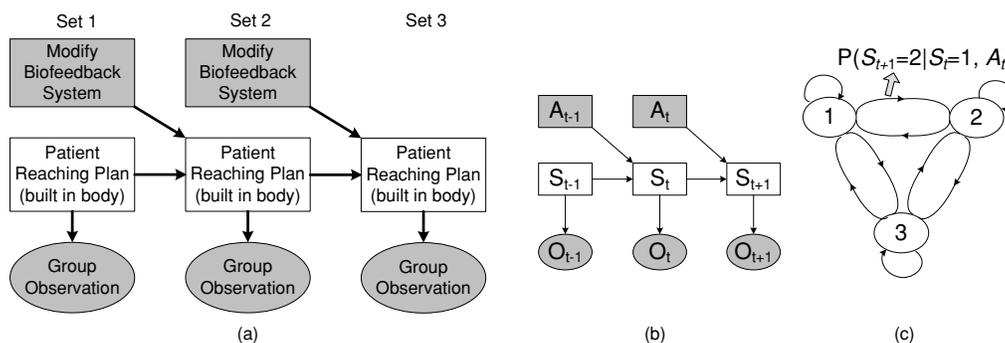


Fig. 10. (a) Graph representation of the rehabilitation procedure. (b) Dynamic Decision Network structure for the media adaptation.  $A_t$  is the decision (i.e., environment change) at time slice  $t$ ,  $S_t$  is the state of the subject movement,  $O_t$  is the observation of the subject movement. Shaded nodes are observable and blank nodes are hidden. Rectangular nodes are discrete and elliptical nodes are continuous. (c) Transition diagram for the hidden node  $S_t$ . The hidden node has three values: 1, 2, and 3. The semantics of the hidden state represent the “goodness” of the body plan to reach the target. The transition probability  $P(S_{t+1}=2|S_t=1, A_t)$  represents the probability of transition from the state 1 to the state 2 under the decision  $A_t$ .

In the following three sections, we discuss the DDN based adaptation models, adaptation recommendation using DDN models and how to use adaptation recommendation in three biofeedback scenarios.

#### 4. DYNAMIC DECISION NETWORK (DDN) BASED ADAPTATION MODEL

In this section, we present the adaptation model based on Dynamic Decision Network (DDN) [Russell and Norvig 2003]. In Section 4.1, we show how a single DDN can be learned and used for the media adaptation. In Section 4.2, we generalize the idea to the mixture of experts.

##### 4.1 Adaptation Model

We use DDN to model the relationship between the adaptation decision and the subject movement. We discuss three key components in the DDN adaptation model: DDN structure, observation prediction and DDN learning in this subsection.

**4.1.1 Dynamic Decision Network (DDN) Structure.** We use the DDN model because it has a similar structure to the rehabilitation procedure (see Section 2.4). Figure 10(a) shows the graph representation

of the rehabilitation procedure. In the first set (at the beginning), the patient has a reaching plan built in his/her body. His/her plan is hidden and is observed by the rehabilitation team when the patient reaches out for ten trials. Then the rehabilitation team has a short discussion and decides how to adapt the biofeedback system in the break. The system adaptation will affect the patient reaching plan for the next set. This procedure repeats until the end of the session. Structurally, the rehabilitation procedure is very similar to the Dynamic Decision Network (DDN) (see Figure 10(b)).

We now propose the general adaptation model in the biofeedback system. Figure 10 (b)–(c) shows the DDN structure for the media adaptation. The time slice in the DDN corresponds to the set (see Section 2.4). This is because we only change the environment (task, audio, visual, physical, and therapist instruction) between the sets and the environment is fixed for all trials within the set. At each time slice  $t$ , there are three nodes: (a) adaptation decision— $A_t$ , (b) state— $S_t$ , and (c) observation— $O_t$ . For the sake of conciseness, we use the term *decision* to refer adaptation decision in rest of this paper. The decision and state nodes are discrete (rectangular nodes in Figure 10(b)) and the observation nodes are continuous (elliptical nodes in Figure 10(b)). The decision and observation nodes are observable (shaded nodes in Figure 10(b)) and the state nodes are hidden (blank nodes in Figure 10(b)). For example, let us assume the DDN model is used for the hand velocity adaptation in the biofeedback system. The adaptation decision  $A_t$  could be the change of musical instrument. The hidden state  $S_t$  represents the goodness of the subject movement plan to achieve smooth velocity. The observation  $O_t$  is the average hand jerkiness over ten trials in the set  $t$ . Hand jerkiness is a well defined measurement in bioengineering. We shall discuss the hand jerkiness later in this paper (Equation (19) in Section 6.3).

We now discuss the nodes in the DDN (see Figure 10(b)) in detail:

- The decision node  $A_t$  represents the media adaptation between the set  $t$  and the set  $t+1$ . The decision is represented by the change of media parameters. For example, assume that we consider the musical instrument and tempo as the media adaptation decision options. Each parameter (instrument or tempo) is a binary variable: (no change/change). Therefore we have four possible decision values by combining these two musical parameters. We make the decision on *changes* rather than on values to make the DDN parameter learning problem tractable.
- The state node  $S_t$  represents the subject hidden state in the reaching/grasping plan acquisition. The semantics of the hidden state represent the “goodness” of the body plan (in movement sequence) to reach the target. We set three possible values for the hidden state which indicate three different states in the reaching/grasping plan acquisition. This is because therapists usually use “good plan”, “bad plan” and “neutral plan” to evaluate the patient reaching movement. Figure 10(c) shows the transition diagram between these three values. The transition probability  $P(S_{t+1}|S_t, A_t)$  is represented as a conditional probability table (CPT):  $R = \{r_{i,j,k}\} = \{P(S_{t+1} = k|S_t = i, A_t = j)\}$ .
- The observation node  $O_t$  represents the subject’s reaching/grasping performance in the set  $t$ . It is the observable output of the hidden state node  $S_t$ . Since the observation node is continuous, we use the probability density  $p$  rather than the probability  $P$ . We use the Gaussian distribution for the sensing probability density  $p(O_t|S_t = i) = N(\mu_i, \Sigma_i)$ , where  $\mu_i$  and  $\Sigma_i$  are the mean vector and covariance matrix when the value of hidden nodes  $S_t$  equals to  $i$ .

**4.1.2 The First-Order Observation Prediction.** In this section, we address the first question posed in Section 3.4. We show how to predict the subject movement performance in the set  $t+1$  (i.e., the *first-order* prediction  $\hat{O}_{t+1}$ ) given the decision and observation sequences for the previous  $t$  sets using the DDN model. Let us denote the decision and observation sequences from the first set to the set  $t$  as  $A_{1:t}$  and  $O_{1:t}$  respectively.

The *first-order* prediction includes three parts: (a) predict the probability of the hidden nodes  $S_{t+1}$ :  $P(S_{t+1}|O_{1:t}, A_{1:t})$ , (b) predict the probability density of the observation node  $O_{t+1}$ :  $p(O_{t+1}|O_{1:t}, A_{1:t})$ , and (c) predict the expected value of the observation node:  $\tilde{O}_{t+1}$ . It is shown in [Russell and Norvig 2003] that  $P(S_{t+1}|O_{1:t}, A_{1:t})$  can be computed as follows:

$$P(S_{t+1}|O_{1:t}, A_{1:t}) = \sum_{S_t} P(S_{t+1}|S_t, A_t) \cdot P(S_t|O_{1:t}, A_{1:t-1}), \quad (1)$$

where  $P(S_t|O_{1:t}, A_{1:t-1})$  can be computed iteratively as follows:

$$P(S_t|O_{1:t}, A_{1:t-1}) = \alpha p(O_t|S_t) \sum_{S_{t-1}} P(S_t|S_{t-1}, A_{t-1}) P(S_{t-1}|O_{1:t-1}, A_{1:t-2}), \quad (2)$$

where  $p(O_t|S_t)$  is the sensing probability density (single Gaussian pdf),  $P(S_t|S_{t-1}, A_{t-1})$  is the transition probability,  $\alpha$  is a normalization constant. Using Equation (1), we can predict the probability density  $p(O_{t+1}|O_{1:t}, A_{1:t})$  as follows:

$$p(O_{t+1}|O_{1:t}, A_{1:t}) = \sum_{S_{t+1}} p(O_{t+1}|S_{t+1}) \cdot P(S_{t+1}|O_{1:t}, A_{1:t}). \quad (3)$$

Therefore, we can compute the conditional mean of  $O_{t+1}$  as the prediction result:

$$\tilde{O}_{t+1}(O_{1:t}, A_{1:t}) = E(O_{t+1}|O_{1:t}, A_{1:t}), \quad (4)$$

where  $E(\cdot)$  is the expectation operator.

**4.1.3 The Second-Order Prediction.** In the similar manner, we can compute the *second-order* prediction  $\tilde{O}_{t+2}(O_{1:t}, A_{1:t+1})$ —(i.e., predict the subject movement performance in the set  $t+2$ ). In order to compute the second order prediction, the rehabilitation team needs to provide two step adaptation decisions after the set  $t$  ( $A_t, A_{t+1}$ ). Note that we do *not* know the observation for the set  $t+1$  (i.e.,  $O_{t+1}$ ). Using Equation (1), we can compute the probability of hidden nodes at the set  $t+2$  (i.e.,  $P(S_{t+2}|O_{1:t}, A_{1:t+1})$ ) as follows:

$$P(S_{t+2}|O_{1:t}, A_{1:t+1}) = \sum_{S_{t+1}} P(S_{t+2}|S_{t+1}, A_{t+1}) \cdot P(S_{t+1}|O_{1:t}, A_{1:t}) \quad (5)$$

where  $P(S_{t+1}|O_{1:t}, A_{1:t})$  is computed using Equation (1). Then we can predict the probability density  $p(O_{t+2}|O_{1:t}, A_{1:t+1})$  as follows:

$$p(O_{t+2}|O_{1:t}, A_{1:t+1}) = \sum_{S_{t+2}} p(O_{t+2}|S_{t+2}) \cdot P(S_{t+2}|O_{1:t}, A_{1:t+1}) \quad (6)$$

The *second-order* prediction is computed as the conditional mean of  $O_{t+2}$  as follows:

$$\tilde{O}_{t+2}(O_{1:t}, A_{1:t+1}) = E(O_{t+2}|O_{1:t}, A_{1:t+1}) \quad (7)$$

where  $E(\cdot)$  is the expectation operator.

**4.1.4 Dynamic Decision Network (DDN) Learning.** We now discuss the learning in the DDN model. There are four kinds of parameters that need to be learned:

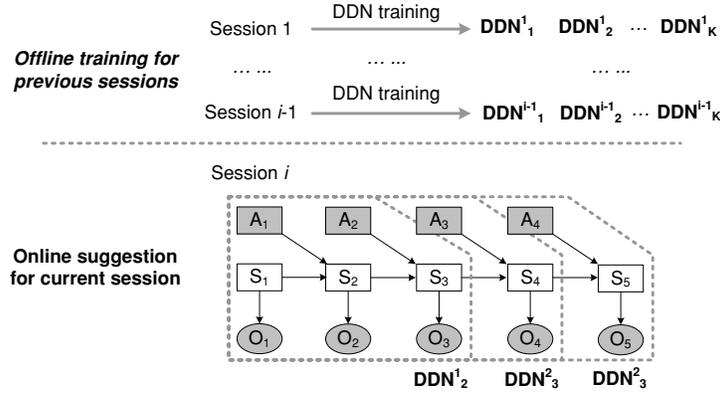


Fig. 11. Diagram for offline training and online adaptation suggestion. We train adaptation models for the session 1 to the session  $i-1$  separately and select a proper model for each set in the session  $i$  to do online suggestion.  $DDN^i_j$  is the  $j$ th DDN model for the  $i$ th session. In this example, we select  $DDN^1_2$ ,  $DDN^2_3$  and  $DDN^3_3$  as the optimal DDNs for the online adaptation suggestion in the set 3, 4, and 5 in the current session.

- (a) Probability of decision:  $\eta_i = P(A_t = i)$ .
- (b) Initial probability of state:  $\pi_i = P(S_1 = i)$ .
- (c) Sensing probability density:  $b_i(o_t) = p(O_t = o_t | S_t = i)$ .
- (d) Transition probability:  $r_{i,j,k} = P(S_{t+1} = k | S_t = i, A_t = j)$ .

We use the EM algorithm [Dempster et al. 1977] to learn these parameters. We assume that we have  $L$  training samples,  $\{y_l\}$ ,  $l = 1, \dots, L$ . Each training sample includes two temporal sequences for decisions and observations respectively ( $y_l = \{a_{1:T_l}^l, o_{1:T_l}^l\}$ ) where  $T_l$  is the length of the sequence for the  $l$ th training sample. For the sake of simplicity, we do not discuss the details of the EM learning in this paper. The details can be found in Dempster et al. [1977]. Note that we assume a uniform prior for the transition probability (i.e.,  $P(S_{t+1} = 1 | S_t, A_t) = P(S_{t+1} = 2 | S_t, A_t) = P(S_{t+1} = 3 | S_t, A_t) = 1/3$ ) if the decision  $A_t$  is not observed in the training samples.

## 4.2 Mixture of Experts

In this section, we extend the basic DDN described in Section 4.1 to a framework using mixture of experts. There are two key ideas:

- (1) In offline training, each session is considered separately. For each session, we use bagging to train  $K$  DDN models.
- (2) In online adaptation, we dynamically search for a proper DDN model from all DDN models trained in previous sessions and use the selected model to give the rehabilitation team adaptation recommendations.

Note that in this paper the offline training and online adaptation is for the same patient. And we do not focus on the adaptation across different patients. This is because the patients have different capabilities and the rehabilitation team has different adaptation strategies for different patients. Figure 11 shows the diagram of offline training and online suggestion. We discuss the offline training in this section and present the online adaptation algorithm in the next section.

**4.2.1 Why Do We Use Mixtures?** We use the mixture of experts (considering each session separately) because we find that the subject performance is weakly correlated. The subject performance is

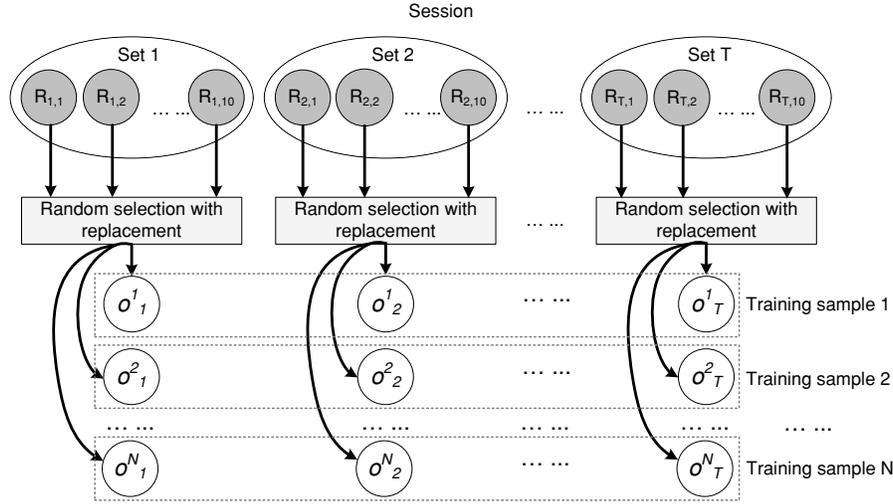


Fig. 12. Construct  $N$  bootstrap training observation sequences.  $R_{i,j}$  is the observation value for the  $j$ th trial in the  $i$ th set.

correlated across different dimensions over different sessions. This prevents a single model to give accurate prediction. We also find that there is *no* strong carryover effect even for two consecutive sessions conducted over consecutive days. By carryover we imply that the performance values of the previous session can *not* be assumed to continue to the current session. This is because subjects may have different amount of use of their impaired arm which results in different physical conditions of their arm ability. Therefore we train different sessions separately. We use bagging [Duda et al. 2000] to train  $K$  DDN models for each session as we have limited training samples.

**4.2.2 Offline Training.** We now present the offline DDN training process using the data from previous sessions. The offline training includes two parts: (a) bootstrap training samples and (b) DDN bagging. For the sake of definiteness, we show the offline training process for one session. The process generalizes for all sessions.

We now show how we generate bootstrap training samples for the DDN learning. Let us assume that we have  $T$  sets for the current session and that each set has ten trials. We compute the observation measurement (e.g., hand jerkiness) for every trial. For each set, we use Huber's method [Huber 1981] to compute the robust mean and variance and detect outliers. Then we randomly select a trial from each set to construct an observation sequence  $o_{1:T}^1$ . We replace this selected trial back to the set and repeat this random selection until we have  $N$  observation sequences—( $o_{1:T}^k, k = 1, \dots, N$ ). For each set, we use uniform distribution on nonoutlier trials and zero probability on outlier trials prior to random selection. Figure 12 shows the diagram for generating  $N$  observation sequences. The decision sequence  $a_{1:T-1}$  includes the decisions taken by the rehabilitation team in the session, where  $a_i$  represents the decision between the set  $i$  and the set  $i+1$ . Therefore, we combine the decision sequence and observation sequence to construct the  $N$  training samples  $y_k = \{a_{1:T-1}, o_{1:T}^k\}$ . Using the EM algorithm (see Section 4.1.4), we can train a DDN model on these  $N$  training samples for the current session.

We repeat the process of generating  $N$  bootstrap training samples  $K$  times and hence create  $K$  training datasets. Each of these bootstrap datasets is used to train a different DDN adaptation model for the current session. Therefore, we train  $K$  DDN models for each session. The final online adaptation recommendation is done by selecting an optimal DDN model from  $K$  DDNs and using its adaptation

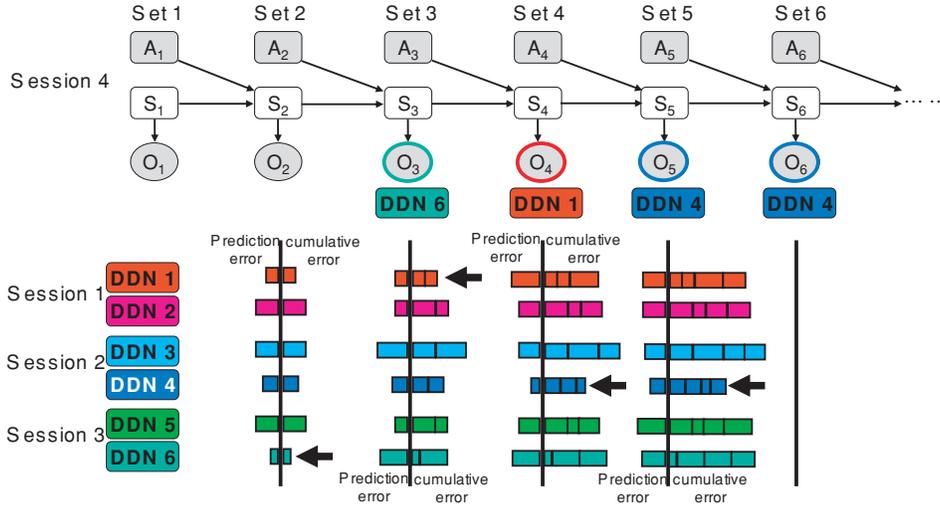


Fig. 13. The performance prediction diagram. In this figure, we assume that we train two DDNs for the first three sessions and use these DDNs to predict performance for the fourth session. At each set in the fourth session, we select the DDN with minimum cumulative error to do prediction.

suggestions. The DDN model is optimal in the sense that it has the closest relationship between the adaptation decision and the movement performance to the current session.

## 5. ADAPTATION RECOMMENDATION

We now present our online adaptation recommendation algorithm based on the mixture of DDN experts to address online adaptation recommendation. We address two questions presented in Section 3.4: (a) performance prediction—given the environment change, we provide the expected subject movement performance for the next set and (b) decision suggestion—given the expectation of subject movement performance, we suggest the optimal change of the environment. We first discuss the performance prediction.

### 5.1 Performance Prediction

The key idea of our performance prediction algorithm is that we first select an optimal DDN from all previous sessions dynamically, and use it to predict the performance for the next set in the current session. The DDN is optimal in the sense that it has the best prediction results (i.e., minimum cumulative prediction error) prior to the current set. Figure 13 shows a performance prediction example.

Let us assume that we have  $Q$  previous sessions. Each session has  $K$  trained DDN adaptation models. Let us denote the  $j$ th DDN trained for the  $i$ th session as  $D_{i,j}$  ( $i = 1, \dots, Q, j = 1, \dots, K$ ). For the current session, we compute the observation for each trial and compute the robust mean [Huber 1981] of the ten trial observations for each set. Let us assume the current session has  $T$  sets. We denote the mean observation of the  $t$ th set as  $o_t$  and the actual decision between the  $t$ th set and the  $t + 1$ th set as  $a_t$ . Therefore the performance prediction is formulated as follows:

1. *The first-order prediction*: predicting the observation for the set  $t + 1$  (i.e.  $o_{t+1}$ ) given  $Q * K$  DDN adaptation models for previous sessions and the decision and mean observation sequences from the first set to the  $t$ th set (i.e.  $a_{1:t}, o_{1:t}$ ).

2. *The second-order prediction*: predicting the observation for the set  $t + 2$  (i.e.  $o_{t+2}$ ) given  $Q^*K$  DDN adaptation models for previous sessions, the mean observation sequence from the first set to the  $t$ th set (i.e.  $o_{1:t}$ ) and the decision sequence from the first set to the  $t + 1$ th set (i.e.  $a_{1:t+1}$ ).

We use the robust mean of the ten trial observations as the set observation because it is an important measurement for therapists to evaluate the subject performance. In this article, we also refer the actual decision taken by the rehabilitation team and the mean observation of set (i.e.,  $a_{1:t}$ ,  $o_{1:t}$ ) as the *ground truth decision* and the *ground truth observation* respectively.

5.1.1 *The First-Order Prediction*. We now discuss how to compute the *first-order* performance prediction. The initial environmental parameters are set by the rehabilitation team based on the observations of physical target reaching at the beginning of the session. At the end of the first set, therapist gives a decision query for the second set (i.e.,  $a_1$ ). Since this is the first prediction for the current session, we do not know which adaptation model can predict well. Hence, we use the median of prediction results of all DDN models as the prediction for the second set:

$$\tilde{o}_2(a_1) = \begin{cases} \text{median}_{D_{i,j} \in \Omega(a_1)}[\tilde{o}(2|D_{i,j}, a_1)] & \text{if } \Omega(a_1) \neq \emptyset \\ \text{median}_{D_{i,j}}[\tilde{o}(2|D_{i,j}, a_1)] & \text{otherwise} \end{cases}, \quad (8)$$

where  $\tilde{o}_2(a_1)$  is the prediction result for the second set for the decision  $a_1$ ,  $\tilde{o}(2|D_{i,j}, a_1)$  is the prediction result for the second set for the decision  $a_1$  by using the DDN model  $D_{i,j}$ ,  $\Omega(a_1)$  is the set of the DDN models  $\{D_{i,j}\}$  that have the decision  $a_1$  in the training samples. The idea is that when we predict the results for the decision  $a_1$ , we should use the previous sessions in which the same decision was taken. If this decision has not been used before (i.e.,  $\Omega(a_1)$  is empty), we assume uniform prior transition probability and to predict the observation. Since we consider all DDN models equivalently, the prediction results may be noisy. Hence, we inform the therapist that this result may not be very reliable.

We can use the previous prediction error to help us find a good DDN adaptation model since the third set. We select the DDN model with minimum cumulative prediction error before the current set as the adaptation model to predict the next set. The *first-order* cumulative prediction error until the current set ( $t$ ) using model  $D_{i,j}$  (i.e.,  $\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t})$ ) is computed as follows:

$$\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t}) = \sum_{k=2}^t \|\tilde{o}(k|D_{i,j}, a_{k-1}) - o_k\|, \quad (9)$$

where  $\tilde{o}(k|D_{i,j}, a_{k-1})$  is the first order prediction result for the  $k$ th set for the decision  $a_{k-1}$  by using the DDN model  $D_{i,j}$  using Equation (4),  $o_k$  is the actual observation for the  $k$ th set. This equation computes the overall prediction error from the beginning of the session to the current set (set  $t$ ) by using the DDN  $D_{i,j}$ . Hence the optimal DDN model  $D^*(t+1|a_t)$  for the prediction of the  $t+1$ th set for the decision  $a_t$  is represented as follows:

$$D^*(t+1|a_t) = \begin{cases} \arg \min_{D_{i,j} \in \Omega(a_t)}[\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t})] & \text{if } \Omega(a_t) \neq \emptyset \\ \arg \min_{D_{i,j}}[\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t})] & \text{otherwise} \end{cases}. \quad (10)$$

Therefore, using Equation (4), we can compute the *first order* prediction results for the  $t+1$ th set by using DDN model  $D^*(t+1|a_t)$ :

$$\tilde{o}_{t+1}(a_t) = \tilde{o}(t+1|D^*(t+1|a_t), a_t), \quad (11)$$

where  $\tilde{o}_{t+1}(a_t)$  is the prediction result for the  $t + 1$ th set for decision  $a_t$ ,  $\tilde{o}(t+1|D^*(t+1|a_t), a_t)$  is the prediction result for the  $t + 1$ th set for the decision  $a_t$  by using the DDN model  $D^*(t+1|a_t)$ .

**5.1.2 The Second-Order Prediction.** In the similar manner, we can compute the *second order* prediction result for the  $t + 2$ th set. The *second-order* cumulative prediction error until the current set ( $t$ ) is computed as follows:

$$\varepsilon_2(t|D_{i,j}, a_{1:t-1}, o_{1:t}) = \sum_{k=3}^t \|\tilde{o}(k|D_{i,j}, a_{k-2}, a_{k-1}) - o_k\|, \quad (12)$$

where  $\tilde{o}(k|D_{i,j}, a_{k-2}, a_{k-1})$  is the second order prediction result for the  $k$ th set for the two step decisions  $a_{k-2}$  and  $a_{k-1}$  by using the DDN model  $D_{i,j}$  using Equation (7),  $o_k$  is the actual observation for the  $k$ th set. In the second order prediction, we select the DDN model with minimum summation of the first order cumulative prediction error and the second order cumulative prediction error as the optimal DDN. The optimal DDN model  $D^*(t + 2|a_t, a_{t+1})$  for the prediction of the  $t + 2$ th set for the adaptation decisions  $a_t, a_{t+1}$  is represented as follows:

$$D^*(t + 2|a_t, a_{t+1}) = \begin{cases} \arg \min_{D_{i,j} \in \Omega(a_t) \cup \Omega(a_{t+1})} [\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t}) + \varepsilon_2(t|D_{i,j}, a_{1:t-1}, o_{1:t})] & \text{if } \Omega(a_t) \cup \Omega(a_{t+1}) \neq \emptyset \\ \arg \min_{D_{i,j}} [\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t}) + \varepsilon_2(t|D_{i,j}, a_{1:t-1}, o_{1:t})] & \text{otherwise} \end{cases} \quad (13)$$

where  $\varepsilon(t|D_{i,j}, a_{1:t-1}, o_{1:t})$  is the first order cumulative prediction error (using Equation (9)) and  $\varepsilon_2(t|D_{i,j}, a_{1:t-1}, o_{1:t})$  is the second order cumulative prediction error (using Equation (12)). Therefore, using Equation (7), we can compute the *second-order* prediction results  $\tilde{o}_{t+2}(a_t, a_{t+1})$  for the  $t+2$ th set by using the DDN model  $D^*(t + 2|a_t, a_{t+1})$ :

$$\tilde{o}_{t+2}(a_t, a_{t+1}) = \tilde{o}(t + 2|D^*(t + 2|a_t, a_{t+1}), a_t, a_{t+1}), \quad (14)$$

where  $\tilde{o}(t+2|D^*(t+2|a_t, a_{t+1}), a_t, a_{t+1})$  is the second-order prediction result for the  $t+2$ th set for the two step decisions  $a_t, a_{t+1}$  by using the DDN model  $D^*(t+2|a_t, a_{t+1})$ .

## 5.2 Decision Suggestion

We now propose the decision suggestion algorithm. The idea is to find the decision whose prediction result is the closest to therapist's expectation (shown in Figure 14). For example, a therapist might say "if I want to improve hand trajectory straightness by 10%, what decision should I take?" In this section, we first discuss a general decision suggestion framework by using the utility function. Then, we apply a simple utility function to provide the online decision suggestion for the rehabilitation team.

The decision suggestion is formulized as follows:

*Online suggestion:* suggest the adaptation decision after the set  $t$ —(i.e.,  $a_t$ ) given  $Q^*K$  DDN adaptation models for previous  $Q$  sessions (train  $K$  DDN models for each session), the decision sequence from the set 1 to the set  $t-1$  (i.e.,  $a_{1:t-1}$ ), the mean observation sequence from the set 1 to the set  $t$  (i.e.,  $o_{1:t}$ ) and the expected performance for the set  $t+1$  (i.e.,  $o_{t+1}^e$ ).

We now discuss how to suggest adaptation decisions based on the utility function. Let us assume that the observation is scale variable in this section. It can be easily extended for the vector variable. Let us assume that there are  $P$  possible decisions in the adaptation model which constructs an decision set  $\Phi = \{\alpha_1, \dots, \alpha_P\}$ . For example, the hand trajectory adaptation scenario has 27 possible decisions

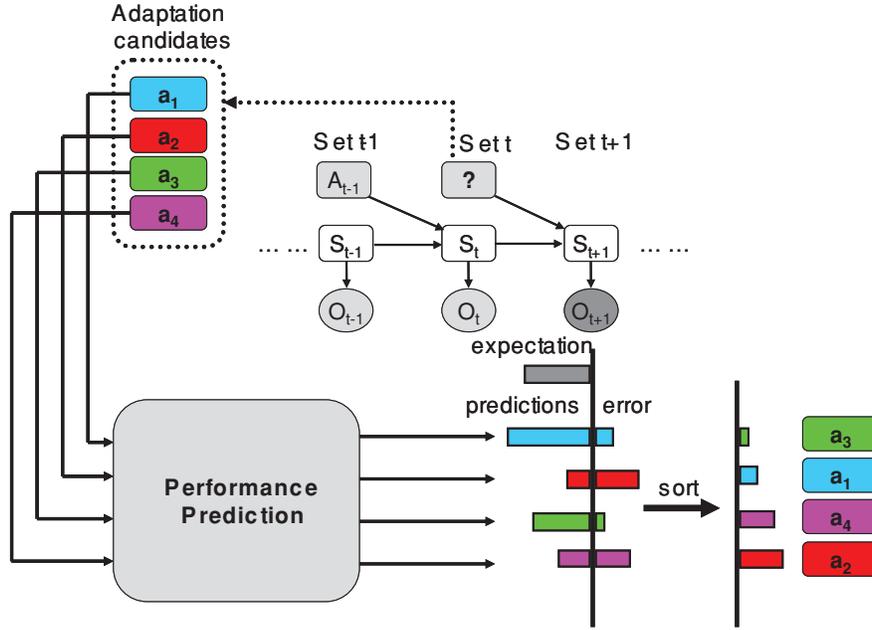


Fig. 14. The adaptation suggestion diagram. In this figure, we assume that we have four possible adaptation decision candidates. For each adaptation candidate, we predict the performance, compare with the expectation from the rehabilitation team and compute the difference between the prediction and the expectation. Finally, we recommend the adaptation candidate with the minimum prediction-expectation difference to the rehabilitation team.

( $P = 27$ ). Let us denote  $\Phi_h$  as a subset of  $\Phi$  in which all decisions are taken in the previous sessions. We only suggest the decisions taken in the previous sessions, because the DDN adaptation models do not learn the cases of decisions that have not been observed and the suggestion may have large error. The optimal decision is the one that minimizes prediction error. This can be achieved by defining an appropriate utility function. Russell and Norvig [2003] show that the optimal decision at time slice  $t$  is the decision that maximizes the expected utility:

$$\begin{aligned} \tilde{a}_t &= \arg \max_{a_t \in \Phi_h} [u_a(a_t)] \\ u_a(a_t) &= \sum_{S_{t+1}} U(S_{t+1}) \left( \sum_{S_t} P(S_{t+1}|S_t, a_t) \cdot P(S_t|o_{1:t}, a_{1:t-1}) \right), \end{aligned} \quad (15)$$

where  $u_a(a_t)$  is the utility of decision  $a_t$ ,  $P(S_{t+1}|S_t, a_t)$  is the transition probability,  $P(S_t|o_{1:t}, a_{1:t-1})$  can be computed using Equation (2),  $U(S_{t+1})$  is the utility function on the state nodes at the time slice  $t+1$ . The utility function can be defined differently for different purposes by the rehabilitation team. In this paper, we define the utility as a function of prediction error:

$$\begin{aligned} U(S_{t+1}|o_{t+1}^e, a_t) &= \lambda_{t+1} (E(O_{t+1}|S_{t+1}) - o_{t+1}^e) \\ \lambda_{t+1} &= \begin{cases} 1 & \text{if } \tilde{o}_{t+1}(a_t) < o_{t+1}^e \\ -1 & \text{otherwise} \end{cases}, \end{aligned} \quad (16)$$

where  $E(O_{t+1}|S_{t+1})$  is the expectation of observation at  $t+1$ th set given  $S_{t+1}$ ,  $o_{t+1}^e$  is the rehabilitation team's expected observation for the  $t+1$ th set,  $\lambda_{t+1}$  is the sign indicator,  $\tilde{o}_{t+1}(a_t)$  is the first order prediction results for the observation at the  $t+1$ th set for the decision  $a_t$  using Equation (11). The utility function is not only related to the state  $S_{t+1}$  but also related to the rehabilitation team's expected observation  $o_{t+1}^e$ , and the decision  $a_t$ . The utility is negative. The absolute value of the utility is the distance from the observation prediction  $\tilde{o}_{t+1}$  to the rehabilitation team's expected observation  $o_{t+1}^e$ . The larger the utility (closer to zero), the smaller the prediction error.  $\lambda_{t+1}$  equals  $+1$  when the predicted observation is less than the rehabilitation team's expected observation and equals  $-1$  otherwise. Maximizing the utility is equivalent to minimizing the prediction error.

In practice, for each possible decision, we search for the DDN adaptation model to maximize the utility (or minimize the prediction error). Then we use the decision which maximizes the utility over all possible decisions as the recommended decision. We represent the decision recommendation in terms of minimizing the prediction error as follows:

$$\tilde{a}_t = \arg \min_{a_t \in \Phi_h} (|\tilde{o}_{t+1}(a_t) - o_{t+1}^e|), \quad (17)$$

where  $\tilde{a}_t$  is the suggested decision between the set  $t$  and the set  $t+1$ ,  $\tilde{o}_{t+1}(a_t)$  is the first order prediction result for the  $t+1$ th set for the decision  $a_t$  (see Equation (11)),  $o_{t+1}^e$  is the therapist's expected observation for the  $t+1$ th set.

## 6. THREE ADAPTATION SCENARIOS IN THE BIOFEEDBACK SYSTEM

We now show how to use the mixture-of-experts based DDN adaptation model in the biofeedback system. We apply the DDN model on three adaptation scenarios: (a) spatial accuracy, (b) straightness of hand trajectory and (c) jerkiness of hand velocity. These three scenarios are important for the stroke patient rehabilitation as they are connected to the *reach*, *open* and *flow* sub-goals respectively (see Section 2.1).

In this article, we consider these three adaptation scenarios separately. Each adaptation scenario has a DDN adaptation model. For these three scenarios, we use the change of media parameter as the decision and we only consider the direction of the change (i.e., increasing/decreasing/no change). The exact quantity of parameter change is decided by the therapist. The domain experts (therapists and artists) suggested that the adaptations should be over three different media spaces: (a) reaching/grasping task control, (b) audio feedback and (c) visual feedback. For each adaptation, therapists determine the goal for the stroke patient rehabilitation.

In the next three sections, we discuss the adaptation scenarios for spatial accuracy, hand trajectory and hand velocity. For each scenario, we discuss the adaptation goal and the decision/hidden state/observation in the DDN model.

### 6.1 Spatial Accuracy Adaptation Scenario

We now propose the adaptation scenario for spatial accuracy. We describe the adaptation goal, the measure of the observation, the meaning of the states and the decisions as follows:

- *Goal*. The goal is to help subject reach close to the target (see Section 3.5).
- *Observation*. We use the distance from the subject hand to the virtual target during grasping as the measurement of spatial accuracy (Figure 15(a)). The virtual target is the center of the grasping zone (Figure 15(a)). It indicates the most comfortable position for the subject to grasp the target. It is calibrated during the calibration when the subject is asked to reach for a physical target comfortably. The subject achieves grasping when his/her hand stay in the grasping zone (Figure 15(a)) for continuous 250ms. The grasping zone is a 3D vertical cylinder around the virtual target. In this paper, we only

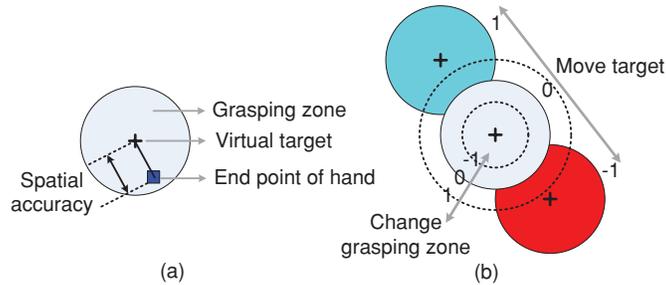


Fig. 15. (a) Diagram for the spatial accuracy, (b) diagram of adaptation decisions for the spatial accuracy scenario.

focus on the grasping zone on the table plane. After the subject successfully grasps the target, the image used in the visual feedback is completed and the sound progression in the audio feedback is completed.

- *State*. The state node has three values (1–3) that indicate three different body computational plans to reach for the target.
- *Decision*. We have nine possible decisions in the adaptation model for the spatial accuracy scenario. These are the combination of the changes of grasping zone radius and virtual target position (see Figure 15(b)). The change of grasping zone has three values (i.e.,  $-1$ : shrink,  $0$ : no change,  $+1$ : enlarge). The change of target position also has three values (i.e.,  $-1$ : move closer,  $0$ : no change,  $+1$ : move further). Note that in the real rehabilitation, not all of these combined decisions happen.

## 6.2 Hand Trajectory Adaptation Scenario

We now present the goal, the measure of the observation, the meaning of the states and the adaptation decisions for hand trajectory adaptation scenario.

- *Goal*. The hand trajectory adaptation is geared towards making the hand trajectory straight (see Section 3.5). Here we focus on the subject's hand trajectory on the horizontal plane (or table plane). Figure 16(a) shows the trajectory along the table plane.
- *Observation*. The straightness of hand trajectory  $G$  is defined as follows:

$$G = \int_0^{\max(z)} |x(z)| dz, \quad (18)$$

where  $x$  and  $z$  are  $x$  and  $z$  coordinates of subject's hand. The  $x$  and  $z$  axis is shown in Figure 16(a). We define the rest position as the origin of the local coordinate system, the direction from the rest position toward the virtual target as the  $z$  axis, the direction that is perpendicular to the  $z$  axis and toward to the left as the  $x$  axis. We represent the  $x$  coordinate as a function of  $z$  coordinate  $-x(z)$ . The straightness  $G$  equals the area of the trajectory region in Figure 16(a). The smaller the area, the straighter the hand trajectory. Through the visual feedback, the subject knows if he/she is out. For example, if the subject's hand moves right further, the image pulls toward to the right (shown in Figure 16 (a)). The visual feedback has a forgiven region, which is named as the  $x$ - $z$  hull. The image pulling effects starts when the subject's hand is close to the boundary of the  $x$ - $z$  hull. The pulling sensitivity increases when the subject hand moves away the hull.

- *State*. In the similar to the spatial accuracy adaptation scenario, the state node of the DDN model in the trajectory adaptation scenario has three values (1–3).

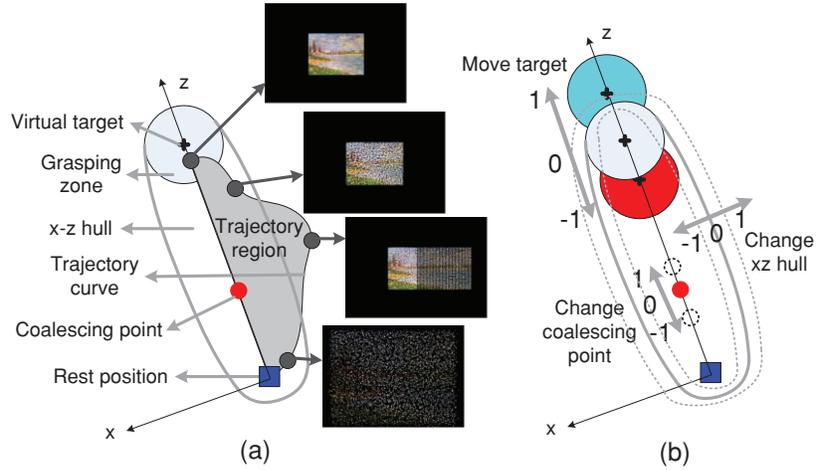


Fig. 16. (a) Trajectory curve and trajectory region along table plane. Four pictures are the visual feedback corresponding to four positions along the trajectory curve. (b) Adaptation decisions for the hand trajectory adaptation scenario (change the x-z hull, change the coalescing point and move the target position).

- *Decision.* We have 27 possible decisions in the hand trajectory adaptation scenario which are the combination of the changes of the x-z hull parameter (size), the coalescing point and the target position (Figure 16(b)). Each change has three values. The x-z hull is a forgiving region for the subject hand movement related to the sensitivity of the image pulling in the visual feedback (Figure 16(a)). The coalescing point is a point on the z axis between the rest position and the virtual target (Figure 16(b)). It indicates the position where the image particles come together (coalescing) and controls the coalescing speed. The position of the virtual target also has effects on the hand trajectory.

### 6.3 Hand Velocity Adaptation Scenario

We now discuss the goal, the observation, the states and the adaptation decisions for the hand velocity adaptation scenario.

- Goal.* The hand velocity adaptation attempts to make the hand velocity smooth (see Section 3.5).
- Observation.* We use a well understood measure (in Bioengineering) for the velocity smoothness (Jerkiness) as the observation of the DDN model for the hand velocity scenario. The jerkiness  $J$  is defined as follows:

$$J = \int_0^T \sqrt{\left(\frac{d^3x}{dt^3}\right)^2 + \left(\frac{d^3y}{dt^3}\right)^2 + \left(\frac{d^3z}{dt^3}\right)^2} dt, \quad (19)$$

where  $T$  is the time when grasping state is achieved,  $x$ ,  $y$  and  $z$  are 3-D coordinates of the subject's hand trajectory. The smaller the value, the smoother the hand velocity.

- State.* In the similar manner, the state node in the hand velocity adaptation DDN has three values (1–3).
- Decision.* We have 12 possible decisions in the hand velocity adaptation. They are related to the changes of musical instrument, tempo, and task difficulty. The change of the musical instrument has

Table I. Number of Sets of All Sessions for Three Subjects

	Session 1	Session 2	Session 3	Session 4	Session 5	Session 6	Total
Subject 1	5	7	4	7	9	9	41
Subject 2	—	7	6	3	7	9	32
Subject 3	—	11	6	10	5	9	41
Total							114

two values (i.e., 1: change, 0: no change). The change of the tempo has three values (i.e.,  $-1$ : decreasing, 0: no change, 1: increasing). The change of the task difficulty has three values ( $-1$ : making the task easier, 1: making the task more challenge, 0: the difficulties of the consecutive sets are in the same level). We select the musical instrument and tempo because the audio feedback Chen et al. [2006] has shown to have a significant effect on the hand velocity. The flow of the subject's hand movement is encouraged by pointillist sound clouds in the main musical line, a smoothly swelling and dipping, wave-shaped musical accompaniment, promotion of synchrony of the involved musical lines and an overall selection of relaxing sound timbres. The hand velocity is also related to the difficulties of the task. When we make the task more challenging, the subject's reaching has more jerkiness. The task difficulty is annotated by the therapist. It relates to multiple media parameters such as the x-z hull, the grasping zone, the target position, etc.

## 7. EXPERIMENTAL RESULTS

We now discuss the experimental results. Our biofeedback system is in a state of continuous improvement. Three stroke patients and twelve nonimpaired subjects are recruited to test the system (30 sessions, 60 hours). Their data are not used in this paper since we keep debugging and improving the system in these 30 sessions.

We recruited another three stroke patients to use our biofeedback system for rehabilitation after the system is finalized. The first patient (subject 1) is a middle-aged female and the second and third patients (subject 2 and subject 3) are aged male. These three patients were suffered mild stroke in the right arm. They were unfamiliar with the system prior to the rehabilitation. Each patient did eight sessions in three consecutive weeks. The eight sessions include one physical pretest, one physical posttest, and six sessions of rehabilitations using our biofeedback system. Each session lasted approximately two hours. The rehabilitations are lead by a physical therapist that has one year experience of using our system.

### 7.1 Experimental Datasets

We now introduce the experimental datasets. Each subject has a dataset. The dataset of the first subject includes six sessions (session 1–6). The second and third subjects have five sessions (session 2–6). The datasets of the second and the third subjects start from the session 2 rather than the session 1 because we found system problems in the first session for both the second and the third subjects. We solved the problems after the first session and the system was not changed from the session 2 to the session 6. The number of sets of all sessions for these three subjects is shown in Table I. We can see there are total 114 sets.

The rehabilitation team did manual adaptations between two consecutive sets within each session for all three subjects. The total number of adaptations for three subjects is 98. As we discussed in the Section 6, the number of possible adaptation decisions for the three scenarios (spatial accuracy, hand trajectory and hand velocity) are 9, 27 and 12 respectively. However, not all possible adaptation decisions are chosen by the rehabilitation team. Table II shows the number of adaptation decisions taken in our experiments.

Table II. Number of Adaptation Decisions Taken in the Experiments for Three Subjects

	Spatial Accuracy	Hand Trajectory	Hand Velocity
Subject 1	5	4	5
Subject 2	6	6	6
Subject 3	5	5	6

## 7.2 Data Sparsity

The adaptation problem needs to be addressed on sparse datasets (i.e., small datasets), out of fundamental characteristics of the problem itself. There are several reasons for the sparsity. The first reason is that the relationship between the adaptation decisions and patient performance is *not* consistent across different patients. This is because patients are very different in terms of their physiological makeup and how the stroke affects their movement. Hence, our rehabilitation team has to develop *different* adaptation strategies for different patients. It is difficult to learn and apply the *same* adaptation model across patients. Therefore, we cannot solve the data sparsity through recruiting many patients at the same time (i.e., learn a single model) since the problem is patient specific. The ideal case is to collect the data for the same patient over a long period (i.e., several years) and then learn an adaptation model tuned to this patient.

Secondly, it is time-consuming and costly to recruit patients over long period. This is because the rehabilitation process is complicated (includes the pretest, the posttest, and the interactive therapy), patients need to come for experiment regularly (two or three times per week), each session takes about two hours and patients need their family members to take them to the experiments (their condition prevents them from driving). Therefore, collecting data for a patient over long period is very difficult.

Finally, the performance of reaching movement is not consistent even within the same patient. This is because his/her physiological conditions are different for different visits (i.e., different sessions). This can happen for many reasons including, taking of medication, and not sleeping correctly (i.e., posture) the night before therapy. This makes the problems more challenging.

## 7.3 Validation Measure

We now propose the validation measure for: (a) performance prediction and (b) decision suggestion. Let us denote  $M$  as number of sessions (e.g.,  $M$  equals 6 for the first subject),  $T_m$  as number of sets for the  $m$ th session,  $a_t^m$  as the ground truth decision (the rehabilitation team's decision) after the  $t$ th set in the  $m$ th session,  $o_t^m$  as the ground truth observation of the  $t$ th set in the  $m$ th session. In this paper, we evaluate our algorithms for three patients separately. The validation measures discussed in this section are in terms of one subject.

**7.3.1 Validation of Performance Prediction.** We use both the mean of relative prediction error ( $\mu_e$ ) and the standard deviation of relative prediction error ( $\sigma_e$ ) to evaluate the performance prediction algorithm. The  $\mu_e$  and  $\sigma_e$  of the *first-order* prediction are computed as follows:

$$\begin{aligned}
 e_t^m &= \frac{|o_t^m - \tilde{o}_t^m(a_{t-1}^m)|}{|o_t^m|} \\
 \mu_e &= \left( \sum_{m=2}^M \sum_{t=2}^{T_m} e_t^m \right) / \left( \sum_{m=2}^M (T_m - 1) \right), \\
 \sigma_e &= \sqrt{\left( \sum_{m=2}^M \sum_{t=2}^{T_m} (e_t^m - \mu_e)^2 \right) / \left( \sum_{m=2}^M (T_m - 1) \right)}
 \end{aligned} \tag{20}$$

where  $\tilde{o}_t^m(a_{t-1}^m)$  is the first order predicted observation of the  $t$ th set in the  $m$ th session given the decision  $a_{t-1}^m$  (see Equation (11)),  $o_t^m$  is the ground truth observation of the  $t$ th set in the  $m$ th session,  $e_t^m$  is the relative prediction error of the  $t$ th set in the  $m$ th session.  $\mu_e$  and  $\sigma_e$  are the mean and standard deviation of relative prediction error over all sets in all sessions.  $(T_m - 1)$  indicates the number of sets for the  $m$ th sessions which we predict the performance. In the similar manner, we can validate the *second-order* performance prediction.

**7.3.2 Validation of Decision Suggestion.** We use three validation measures to evaluate the decision suggestion: (a) average rank, (b) relative utility ratio (RUR) and (c) difference of prediction error (DPE). As we discussed in the Section 5.2, we need to know the rehabilitation team's expected observation before the decision suggestion. We use the actual observation  $o_{t+1}^m$  (i.e., observation of the subject performance (see Section 5.1)) as the expected observation and use the actual decision  $a_t^m$  (i.e., the decision made by the rehabilitation team) as the ground truth decision. This is because the observation  $o_{t+1}^m$  is actually caused by the actual decision  $a_t^m$ .

The *average rank* is the average rank of the ground truth decision (made by the rehabilitation team) over all sets in all sessions. As we discussed in the Section 5.2, we sort all possible decisions by the utility (see Equation (15)) in the descending order and select the first rank decision. We use average rank of the ground truth decision to indicate the accuracy of our decision suggestion algorithm. If the ground truth decision has low rank, it has the large utility to achieve the rehabilitation team's expectation for the next set. The lower the average rank, the more accurate the decision suggestion. The ideal case is that the average rank equals to one. This means that all ground truth decisions are recommended. The average rank of the ground truth decision for the *first-order* decision suggestion is computed as follows:

$$r_{ave} = \left( \sum_{m=2}^M \sum_{t=2}^{T_m} rank(a_t^m | o_{t+1}^m) \right) / \left( \sum_{m=2}^M (T_m - 1) \right), \quad (21)$$

where  $rank(a_t^m | o_{t+1}^m)$  is the rank of the ground truth decision  $a_t^m$  over all possible decisions given the rehabilitation team's expected observation  $o_{t+1}^m$ .

However, the average rank does *not* consider the utility difference between the ground truth decision and our recommendation. For example, if the ground truth decision has higher rank, its utility may be very close to the first rank decision. We should consider it differently with other high rank decisions with much smaller utility. Therefore, we use the relative utility ratio (RUR)  $r_u$  to evaluate the online suggestion algorithm. The RUR of the decision suggestion is computed as follows:

$$r_u(\{a_t^m\}) = \frac{E_t[u_a(\hat{a}_t^m(1))] - E_t[u_a(a_t^m)]}{u_{\max} - E_t[u_a(\hat{a}_t^m(1))]}, \quad (22)$$

where  $\{a_t^m\}$  is the set of the ground truth decisions,  $u_a(\hat{a}_t^m(1))$  is the utility (Equation (15)) of the first rank decision after the  $t$ th set in the  $m$ th session,  $u_a(a_t^m)$  is the utility of the ground truth decision  $a_t^m$ ,  $u_{\max}$  is the maximum of possible utility and  $E_t$  is the expectation operator on set. In this paper, the maximum utility  $u_{\max}$  is zero which means the predicted observation is exactly the same as the rehabilitation team's expected observation. For the ideal case in which all ground truth decisions are selected as the first rank decision, the relative utility ratio equals to zero. The smaller the relative utility ratio, the more accurate the online suggestion.

The difference of prediction error (DPE) measures the difference between the ground truth decision and our suggested decision (the first rank decision) in terms of the prediction error. The difference of

prediction error (DPE) for the decision suggestion is computed as follows:

$$d(\{a_t^m\}) = E_t \left( \frac{|o_t^m - \tilde{o}_t^m(a_{t-1}^m)| - |o_t^m - \tilde{o}_t^m(\hat{a}_{t-1}^m(1))|}{|o_t^m|} \right), \quad (23)$$

where  $o_t^m$  is the ground truth observation for the  $t$ th set in the  $m$ th session,  $\tilde{o}_t^m(a_{t-1}^m)$  is the predicted observation of the  $t$ th set in the  $m$ th session given the ground truth decision  $a_{t-1}^m$  (see Equation (11)),  $\tilde{o}_t^m(\hat{a}_{t-1}^m(1))$  is the predicted observation of the  $t$ th set in the  $m$ th session given the recommended decision (i.e., first rank decision  $\hat{a}_{t-1}^m(1)$ ) and  $E_t$  is the expectation operator on set.

#### 7.4 Results of Performance Prediction

In this section, we discuss the experimental results of the performance prediction. We first explain the training and testing processes. Secondly, we present the first-order performance prediction results for three stroke patients for three adaptation scenarios. Then we compare the prediction results between the single model and the mixture of experts. Finally, we show the *second-order* performance prediction results.

**7.4.1 Training and Testing.** We now introduce the training and testing process. We consider three subjects separately for both training and testing. This means that we first do *not* train Dynamic Decision Network (DDN) models across patients. Secondly, we do *not* use the DDN models trained from other patients to do prediction for the current patient. This is because the patients have different capabilities, and our adaptation strategies in the rehabilitations vary for different patients. Therefore, for each patient, we use his/her rehabilitation history (i.e., previous sessions) to train the DDN models and apply these models to do the performance prediction for his/her current session. For each patient, we apply training and testing for three adaptation scenarios. They are: (a) adaptation for spatial accuracy, (b) adaptation for hand trajectory and (c) adaptation for hand velocity.

*Training.* For every subject, for each scenario, we train  $K = 10$  (see Section 4.2.2) different DDN models for each session. For each DDN model training, we generate  $N = 50$  observation training sequences using bootstrap technique (see Section 4.2.2).

*Testing.* the ground truth observation for each set is the robust mean observation over ten trials in the set. The ground truth adaptation decision is made by the rehabilitation team. In this section, we show the prediction results for both online prediction and offline prediction. We explain the online prediction and the offline prediction as follows:

1. *Online prediction.* We predict the performance for the current session using the DDN models trained from all previous sessions for the same patient. For example, for the subject 2, we predict the observation for the session 5 using the DDN models for the session 2, 3, 4 (the session 1 is removed due to system problems).
2. *Offline prediction.* We predict the performance for a session using all previous sessions and succeeding sessions for the same patient. For example, for the subject 2, we predict the observation for the session 5 using the DDN models for the session 2, 3, 4, 6. While in practice, only online predictions are possible. The offline results are indication of the possible improvement of the results with more data.

**7.4.2 The First-Order Performance Prediction Results.** We now present the *first-order* performance prediction results. Figure 17 shows the *first-order* prediction results for three stroke patients for three adaptation scenarios: spatial accuracy (left column), hand trajectory (middle column) and hand velocity (right column). Note, we do not show the prediction results of the first session for the subject 1 and the prediction results of the second session for the subject 2 and 3. This is because that they are the beginning of rehabilitation and we do not have previous rehabilitation data to learn models.

Table III.

The *first-order* online/offline prediction error for three subjects for three adaptation scenarios. The mean  $\mu_e$  and std  $\sigma_e$  of relative prediction error are computed using Equation (20)

		Spatial accuracy		Hand trajectory		Hand velocity	
		mean $\mu_e$	std $\sigma_e$	mean $\mu_e$	std $\sigma_e$	mean $\mu_e$	std $\sigma_e$
Subject 1	online	8.58%	7.97%	17.85%	17.66%	13.76%	19.17%
	offline	6.02%	6.86%	12.64%	13.40%	8.00%	9.08%
Subject 2	online	8.97%	8.10%	19.48%	19.96%	14.49%	11.77%
	offline	7.25%	6.84%	16.17%	18.94%	11.04%	9.20%
Subject 3	online	19.36%	17.90%	14.86%	18.07%	16.30%	23.68%
	offline	14.11%	13.59%	9.23%	12.20%	9.82%	15.92%

Table III shows the mean and standard deviation of relative prediction error over three patients using Equation (20).

We have two observations from the experimental results:

1. *Online prediction.* The online prediction errors for three adaptation scenarios (spatial accuracy, hand trajectory and hand velocity) for all three patients are low. The mean relative prediction errors are less than 20%, and the standard deviations of relative prediction error are less than 24%. This indicates that our DDN model works well to model the relationship between the patient movement and the media adaptation. We also observe the prediction error at the beginning of session is larger (especially for the second set) and the error decreases as the set number increases (see Figure 17). This is because our expert selection is based on the prediction error of previous sets. The probability to find the proper DDN expert increases as more sets come. We also notice that the predictions are *not* very good for some sessions (such as session 2, 4 in the spatial accuracy, the session 3, 5 in the hand trajectory, session 4 in the hand velocity for subject 1, session 3 in three adaptation scenarios for subject 2, session 4 in the spatial accuracy, session 3, 5 in the hand velocity for subject 3). This is because the movement-adaptation relationship in previous sessions does not repeat in the current session.
2. *Comparison between the online prediction and the offline prediction.* The offline prediction has smaller prediction errors for sessions except the last session for all three subjects (The online prediction is equivalent to the offline prediction for the last session). This is because that we can only use the early session to predict the late session in the online prediction. While in the offline prediction, the similarity between the early session and the late session in terms of the movement-adaptation relationship results in good predictions for both the early session and the later session. For example, we can see that for the subject 1, the prediction results of the session 2 and 4 for the spatial accuracy, the session 2 for the hand trajectory and the session 4 for the hand velocity are improved significantly by using the offline prediction. We understand that in practice, we can only do online prediction. We use the offline prediction to show that the online prediction results for some early sessions are not good because we do not have enough data. As we get more data, our algorithm provides more accurate prediction. In average over the three subjects over the three adaptation scenarios, the offline prediction reduces the mean and standard deviation by 30.0% and 25.3% respectively compared to the online prediction. Note that in the rest of this paper, when we mention “in average” in the result comparison, it refers to in average over the three subjects over the three adaptation scenarios.

**7.4.3 The Comparison between the Single Model and the Mixture of Experts.** We now compare the online prediction results between using the single model and the mixture of experts. First, let us compare the training and testing processes between the single model and the mixture of experts. Here we focus

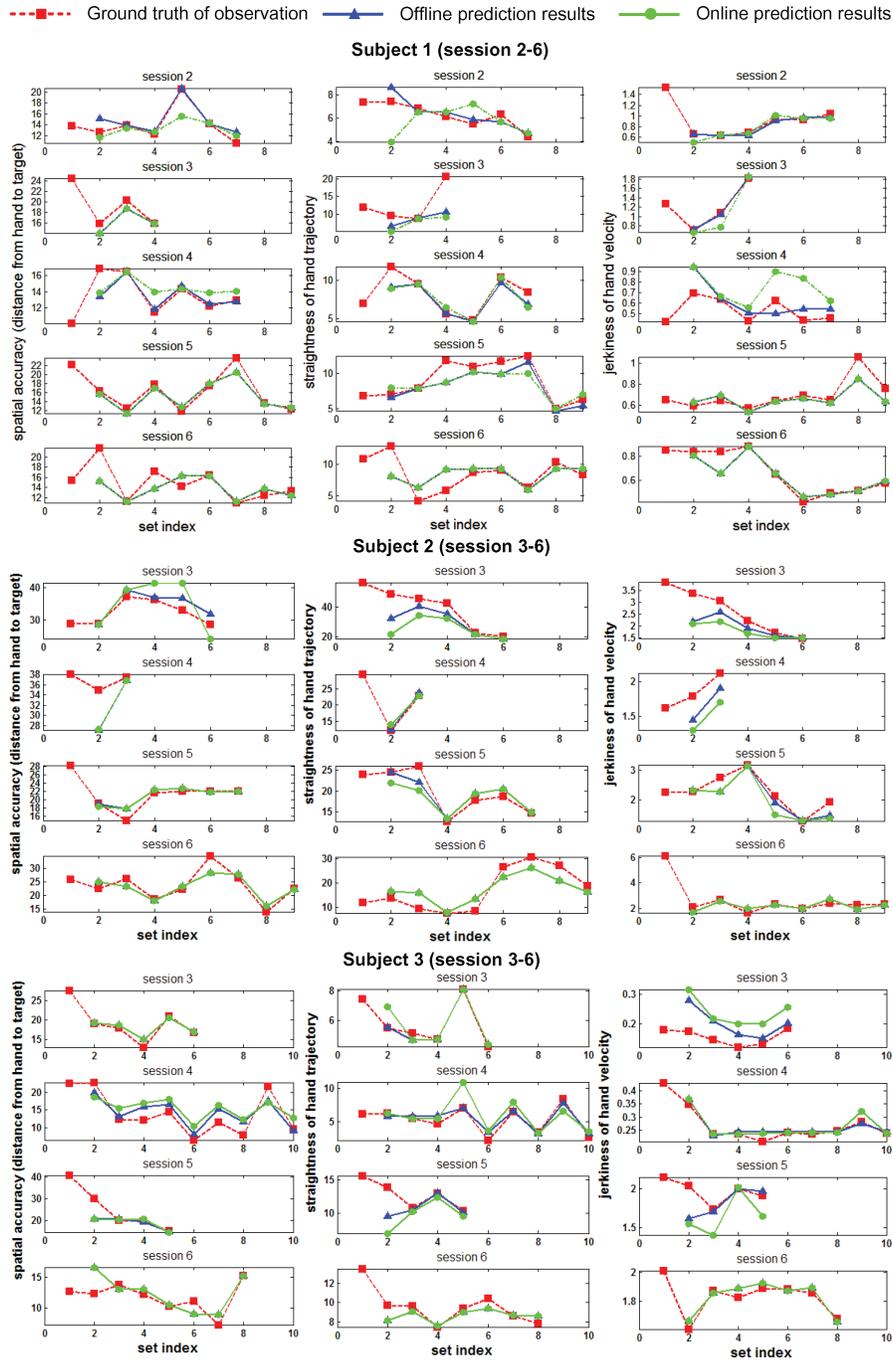


Fig. 17. Experimental results for the *first-order* performance prediction for three stroke patients. For each subject, the left column is the prediction results of the spatial accuracy adaptation. The middle column is the prediction results of the hand trajectory adaptation. The right column is the prediction results of the hand velocity adaptation.

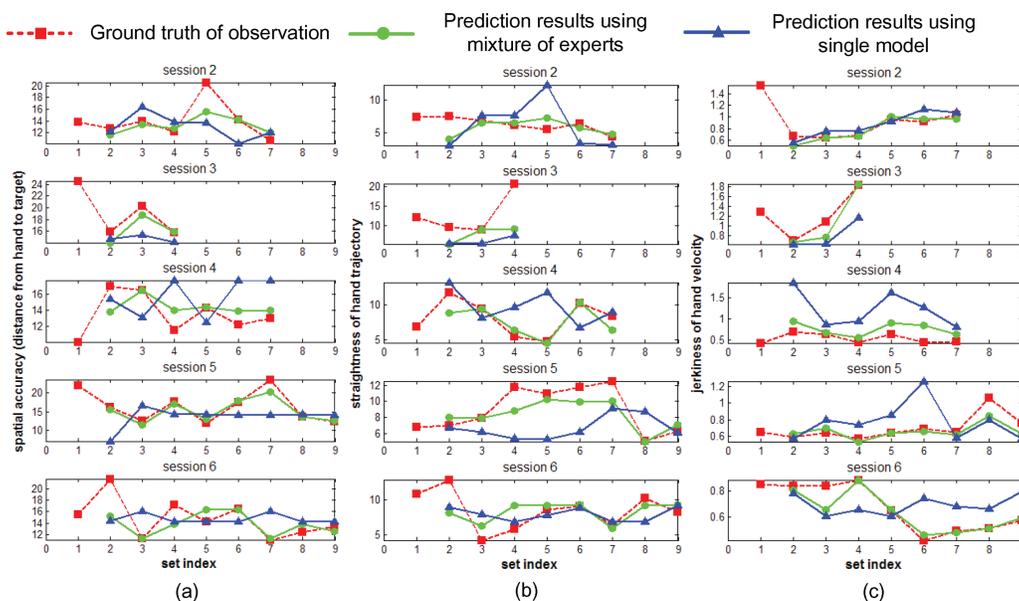


Fig. 18. The comparison of prediction error (the first order online prediction) between the single model and the mixture of experts for the first stroke patient for three adaptation scenarios. The left column is the prediction results for the spatial accuracy scenario, the middle column is the prediction results for the hand trajectory scenario, and the right column is the prediction results for the hand velocity scenario.

on the training and testing for one subject for one scenario. In the similar manner, we can train/test for other subjects and other scenarios. Let us assume that we do spatial accuracy adaptation for the subject 1, and the current session is session  $k$ .

#### —Single Model

—*Training*: we train one DDN model for all previous sessions (session 1 –  $k-1$ ). For each previous sessions, we generate  $N = 50$  observation training sequences using bootstrap technique. Thus, we have  $N(k-1)$  training samples to train one DDN model.

—*Testing*: use the trained DDN model to do performance prediction for the current session (session  $k$ ).

#### —Mixture of Experts

—*Training*: we train  $K = 10$  different DDN models for each previous session (session 1 –  $k-1$ ). For each DDN model training, we generate  $N = 50$  observation training sequences using bootstrap technique. Thus, we have  $NK(k-1)$  training samples to train  $K(k-1)$  DDN models and each model have  $N$  training samples.

—*Testing*: we select a proper DDN model to predict the performance for every set in the current session.

Figure 18 shows the online prediction results of the first subject using both the single model and the mixture of experts. We can see the prediction results using the mixture of expert are much closer to the ground truth. Table IV shows the comparison of relative prediction error between using the single model and using the mixture of experts for three stroke patients for three adaptation scenarios. We observe that the mixture of experts significantly reduces the prediction error compared to the single model. In average over three patients over three adaptation scenarios, using the mixture of experts reduces the

Table IV.

The comparison of prediction error (the first order online prediction) between the single model and the mixture of experts for three stroke patients for three adaptation scenarios. The mean and standard deviation of the relative prediction error are computed using Equation (20)

		Spatial accuracy		Hand trajectory		Hand velocity	
		mean	std	mean	std	mean	std
Subject 1	single	23.00%	15.22%	39.35%	34.62%	44.65%	49.80%
	mixture	8.58%	7.97%	17.85%	17.66%	13.76%	19.17%
Subject 2	single	19.78%	16.30%	38.12%	24.11%	36.42%	30.05%
	mixture	8.97%	8.10%	19.48%	19.96%	14.49%	11.77%
Subject 3	single	42.90%	45.30%	31.53%	32.31%	33.94%	35.58%
	mixture	19.36%	17.90%	14.86%	18.07%	16.30%	23.68%

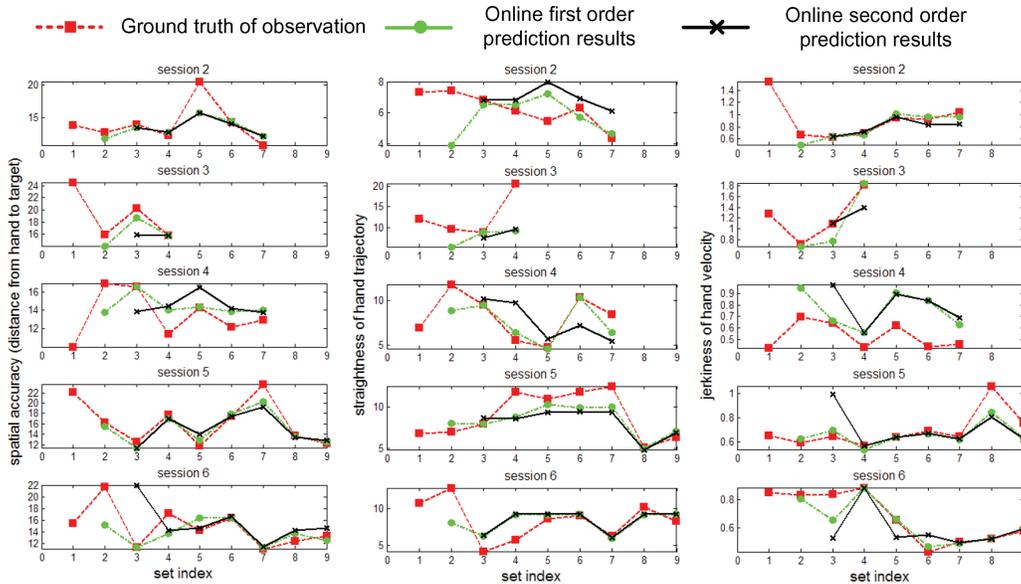


Fig. 19. Experimental results of the *second-order online* performance prediction for the first subject for three adaptation scenarios. Left column: Prediction results of the spatial accuracy adaptation. Middle column: Prediction results of the hand trajectory adaptation. Right column: Prediction results of the hand velocity adaptation.

mean and standard deviation of prediction error by 56.7% and 47.4% respectively compared to using the single model. These results prove that our mixture-of-experts framework works very well.

**7.4.4 The Second-Order Performance Prediction Results.** We now present the experimental results for the *second-order* performance prediction (see Equation (14) in Section 5.1.2). Figure 19 shows the results of the *second-order online* performance prediction for the first subject for three adaptation scenarios. Figure 20 shows the results of *second-order offline* performance prediction for the first subject for three adaptation scenarios. Table V shows the prediction error of the *second-order* prediction for three stroke patients for three scenarios. We have three observations:

1. *Online prediction.* The online prediction errors for the *second-order* prediction for three stroke patients for three adaptation scenarios are low. The mean are less than 26%, and the standard deviation are less than 33.5%. This indicates that our mixture-of-experts based DDN model works well to model

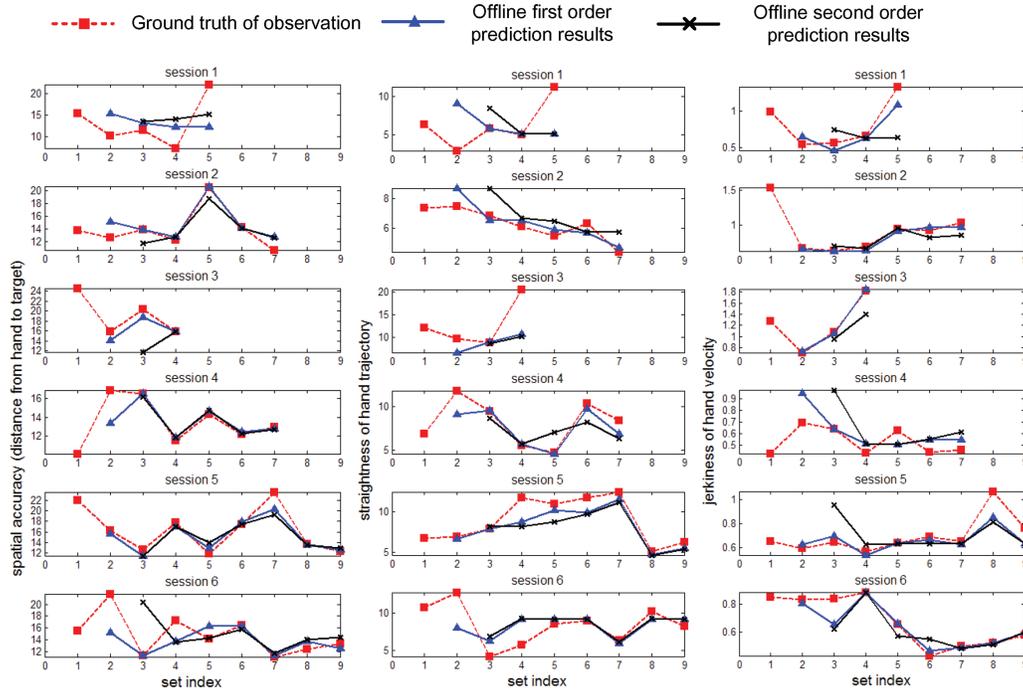


Fig. 20. Experimental results of the *second-order offline* performance prediction. Left column: Prediction results of the spatial accuracy adaptation. Middle column: Prediction results of the hand trajectory adaptation. Right column: Prediction results of the hand velocity adaptation.

Table V.

The *second-order* online/offline prediction error for three subjects for three adaptation scenarios. The mean  $\mu_e$  and std  $\sigma_e$  of relative prediction error are computed using Equation (20)

		Spatial accuracy		Hand trajectory		Hand velocity	
		mean	std	mean	std	mean	std
Subject 1	online	13.42%	18.47%	23.14%	20.48%	19.94%	23.25%
	offline	11.11%	16.95%	19.82%	17.90%	15.99%	13.90%
Subject 2	online	13.78%	18.34%	24.51%	28.68%	19.96%	16.33%
	offline	11.51%	15.48%	22.41%	27.88%	15.29%	12.28%
Subject 3	online	24.65%	31.93%	16.52%	20.30%	25.97%	33.36%
	offline	18.24%	23.00%	14.49%	16.12%	22.56%	23.45%

the relationship between patient movement and media adaptation not only in the first order prediction but also in the second order prediction.

2. *Comparison between the first-order prediction and the second-order prediction.* The second-order prediction introduces more prediction errors compared to the first-order prediction. This is because the second-order prediction has less observation information than the first-order prediction. For the online prediction, the mean and standard deviation of prediction error of the second-order prediction increase by 38.3% and 55.8% respectively in average compared to the first-order prediction. For the offline prediction, the mean and standard deviation of prediction error increases by 67.0% and 65.1% in average from the first-order prediction to the second-order prediction. These numbers are large

Table VI.

The average rank of decision suggestion for three stroke patients for three adaptation scenarios. The average rank is computed by using Equation (21). The average rank for the spatial accuracy for the first subject 1.65/5 means the average rank is 1.65 and the number of observed decisions in rehabilitation is 5.

Scenarios \ Subject	Spatial accuracy	Hand trajectory	Hand Velocity
Subject 1	1.65/5	1.48/4	1.77/5
Subject 2	1.90/6	1.71/6	1.57/6
Subject 3	1.72/5	1.56/5	1.72/6

because the prediction errors of the first-order prediction for both mean and standard deviation are small (i.e., small denominators). However, in terms of the absolute difference of prediction errors between the first-order prediction and the second-order prediction, the mean and standard deviation increase about 5.4% and 7.4% of the ground truth value in average for the online prediction and increase about 6.4% and 6.8% of the ground truth value in average for the offline prediction. These low absolute difference of prediction errors (5%–8%) indicate that the mixture-of-experts based DDN model works well for higher order prediction. Although the second-order prediction has more prediction errors, the prediction errors are still low. Therefore, the second-order prediction can provide useful predictions for the next two sets for the rehabilitation team. This is very important for the rehabilitation team to make the decisions for the next two steps rather than one.

3. *Comparison between the online prediction and the offline prediction for the second-order prediction.* The offline prediction has smaller prediction errors than the online prediction for the case of second-order prediction. The offline prediction reduces the mean and standard deviation of relative prediction error by 16.8% and 20.6% in average respectively compared to the online prediction. This is also because the online prediction results for some early sessions are not good because we do not have enough data. As we get more data, our algorithm provides more accurate prediction.

## 7.5 Results of Decision Suggestion

In this section, we discuss the experimental results for the online decision suggestion. We first discuss the decision suggestion results using the mixture-of-experts based DDN model for three stroke patients for three adaptation scenarios. Then we compare the decision suggestion results between using the single model and using the mixture of experts.

**7.5.1 The Decision Suggestion Results Using the Mixture of Experts.** We now present the decision suggestion results using the mixture-of-experts based DDN model for three stroke patients for three adaptation scenarios: spatial accuracy, hand trajectory and hand velocity. Table VI shows the average rank (see Equation (21)) of the ground truth decisions. We can see that the average rank is close to one. This indicates that the ground truth decisions have larger utility compared with other decision options. There are two reasons why some ground truth decisions are not chosen as the first rank decision. First, the performance predictions (related to the utility) for the ground truth decisions are not accurate for some sets. Second, other decisions may result in smaller prediction error compared to the ground truth decisions.

Table VII shows the relative utility ratios (RURs) (see Equation (22)) of the decision suggestion for three stroke patients for three adaptation scenarios. We observe that the relative utility ratios (RURs) of the ground truth decisions are close to zero in most cases. This indicates that our online suggestion algorithm works well. For a few cases (e.g., the hand velocity scenario for the subject 1 and the spatial accuracy scenario for the subject 2), the RURs are a little larger. But we find that the differences of

Table VII.

The relative utility ratios (RURs) (see Equation (22)) of decision suggestion for three subjects for three adaptation scenarios. For each subject, we show the results for three adaptation scenarios (spatial accuracy, hand trajectory and hand velocity). For each scenario, the first column (GT) shows the relative utility ratios (RURs) of the ground truth decisions, the second column (second) shows the RURs of the second rank decisions and the last column shows the RURs of the last rank decisions.

Subject	Spatial accuracy			Hand trajectory			Hand Velocity		
	GT	second	last	GT	second	last	GT	second	last
Subject 1	0.47	1.31	6.44	0.12	0.99	2.76	0.92	1.85	6.86
Subject 2	0.78	1.22	6.11	0.29	0.75	5.06	0.38	1.41	7.62
Subject 3	0.56	0.62	4.19	0.52	1.99	7.99	0.34	1.36	20.68

Table VIII.

The differences of prediction error (DPEs) (see Equations (23)) of decision suggestion for three subjects for three adaptation scenarios. For each subject, we show the results for three adaptation scenarios (spatial accuracy, hand trajectory and hand velocity). For each scenario, the first column (GT) shows the DPEs of the ground truth decisions, the second column (second) shows the DPEs of the second rank decisions, and the last column shows the DPEs of the last rank decisions.

Subject	Spatial accuracy (%)			Hand trajectory (%)			Hand Velocity (%)		
	GT	2nd	last	GT	2nd	last	GT	2nd	last
Subject 1	1.78	7.02	35.63	1.54	16.83	44.59	5.79	11.19	38.35
Subject 2	3.41	6.27	30.07	5.91	8.47	66.71	3.95	12.79	74.74
Subject 3	5.49	9.17	44.24	3.57	19.37	76.16	2.36	4.10	215.31

prediction error (DPEs) (see Equation (23)) of these cases are very small (see Table VIII). This means that although the prediction error of the ground truth decisions is very close to the first rank suggestions, the prediction error of the first rank decisions is close to zero which results in a small denominator in Equation (22). We also compare the relative utility ratios (RURs) of the ground truth decisions to the second rank decisions and last rank decisions. The RURs for the second rank decisions can be computed by using the second rank decisions as the input of Equation (22) and Equation (23) (i.e.,  $r_u(\{\hat{a}_i^m(2)\})$ ). In the similar manner, we can compute the RURs for the last rank decisions. We can see that the RURs of the ground truth decisions are significantly less than the second rank decisions and the last rank decisions. This indicates that although some ground truth decisions are not be suggested as the first option, they are very close to the first rank suggestions in terms of prediction error.

Table VIII shows the differences of prediction error (DPEs) (see Equation (23)) of the decision suggestion for three stroke patients for three adaptation scenarios. We observe that the differences of prediction errors (DPEs) of the ground truth decisions are very close to zero for all cases. The small DPEs mean that although some ground truth decisions are not rank first, their prediction errors are very close to the first-rank suggestions. We compare the DPEs for the ground truth decisions to the second-rank decisions and last rank decisions. The DPEs for the second rank decisions can be computed by using the second rank decisions as the input of the Equation (22) and Equation (23) (i.e.,  $d_u(\{\hat{a}_i^m(2)\})$ ). In the similar manner, we can compute the DPEs for the last-rank decisions. We can see that the DPEs for the ground truth decisions are significantly less than the second-rank decisions and the last-rank decisions. This indicates that our online suggestion algorithm works well.

**7.5.2 The Comparison between the Single Model and the Mixture of Experts.** We now compare the decision suggestion results between using the single model and the mixture of experts. The training

Table IX.

The comparison of decision suggestion results between using the single model and the mixture of experts for three subjects for three adaptation scenarios. For each subject, we show the results for three adaptation scenarios (spatial accuracy, hand trajectory and hand velocity). For each scenario, we compare the average rank, the relative utility ratio (RUR) (see Equation (22)) and the difference of prediction error (DPE) (see Equation (23)) between using the single model and the mixture of experts.

Subject	model	Spatial accuracy			Hand trajectory			Hand Velocity		
		rank	RUR	DPE (%)	rank	RUR	DPE (%)	rank	RUR	DPE (%)
1	single	2.35	1.31	13.26	1.84	0.96	20.17	2.29	1.32	26.02
	mixture	1.65	0.47	1.78	1.48	0.12	1.54	1.77	0.92	5.79
2	single	2.81	1.27	19.73	2.92	1.32	45.77	1.85	0.44	10.96
	mixture	1.90	0.78	3.41	1.71	0.29	5.91	1.57	0.38	3.95
3	single	2.11	1.13	18.87	2.11	1.23	27.18	2.09	1.46	31.27
	mixture	1.72	0.56	5.49	1.56	0.52	3.57	1.72	0.34	2.36

and testing processes for the single model and the mixture of experts are introduced in Section 7.4.3. Table IX shows the comparison of decision suggestion results (average rank, relative utility ratio (RUR) and difference of prediction error (DPE)) between using the single model and using the mixture of experts for three stroke patients for three adaptation scenarios. We observe that the mixture of experts significantly reduces the average rank, RUR and DPE compared to the single model. In average, using the mixture of experts reduces the average rank, RUR and DPE by 24.5%, 55.9%, and 82.5% respectively compared to using the single model. These results prove that our mixture-of-experts framework works very well.

## 8. CONCLUSION

This paper presented a new framework for media adaptation, for biofeedback rehabilitation. The problem was challenging for several reasons—(a) high dimensionality of parameter space, (b) within session and across session subject performance variation and (c) domain expert decision making was a non first order Markov process. Our key insight is to understand media adaptation as a real-time feedback control problem. The determination of the duality between the mediated rehabilitation problem and the robot navigation problem allowed us to leverage mathematical formalisms that have worked well in other domains.

We used a mixture-of-experts based Dynamic Decision Network (DDN) for online media adaptation. The mixture of experts model was adopted after we realized that patients exhibit significant variability across sessions. The expert mixture was trained using the familiar EM algorithm. The models were used to answer two basic questions—(a) given a specific adaptation suggested by the domain expert, predict expected patient performance and (b) given the expected performance, determine optimal media adaptation decision. The questions are answered through an optimality criterion based search on DDN models trained in previous sessions.

We used the Dynamic Decision Network (DDN) to model the relationship between the media adaptation and the patient performance. The decision network contains three types of nodes—adaptation decision ( $A_t$ , observable), state describing the movement plan ( $S_t$ , hidden) and subject movement observations ( $O_t$ , observable). The decision node  $A_t$  represents the media adaptation decision after subject movement observations at time  $t$ . We make decision on *changes* rather than on values to make the DDN parameter learning problem tractable. The state node  $S_t$  represents the subject hidden state in reaching/grasping plan acquisition. The semantics of the hidden state represent the “goodness” of the body plan (in movement sequence) to reach the target. The observation node  $O_t$  represents the subject’s reaching/grasping performance in the set  $t$ . We train DDN mixtures per patient, per session. In online

adaptation, we dynamically search for a proper DDN model from all DDN models trained in previous sessions and use the selected model to recommend adaptations for the rehabilitation team.

Our experiments on the real stroke patient data show excellent results on both performance prediction and adaptation decision recommendation for three stroke patients for three adaptation scenarios: special accuracy, hand trajectory and hand velocity. Compared to the single model, the mixture of experts significantly improves both the performance prediction and the decision suggestion. We plan to extend our research in several ways—(a) extend decision making to changes in the physical setup, and suggestions on when the therapist should discuss with the patient. (b) joint decision making by incorporating the relationship between different aspects of patient’s movements (e.g., combine spatial accuracy, hand trajectory and hand velocity). We also plan to continue our patient rehabilitation trials over the next few months.

## REFERENCES

- ABOWD, G. D., MYNATT, E. D., AND RODDEN, T. 2002. The human experience [of ubiquitous computing]. *IEEE Pervasive Comput.* 1, 48–57.
- BROOKS, R. A. 1991a. Intelligence Without Reason. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 569–595.
- BROOKS, R. A. 1991b. Intelligence without representation. 47, 139–159.
- BROOKS, R. A., COEN, M., DANG, D., DEBONET, J., KRAMER, J., LOZANO-PEREZ, T., MELLOR, J., POOK, P., STAUFFER, C., STEIN, L., TORRANCE, M., AND WESSLER, M. 1997. The intelligent room project. In *Proceedings of the 2nd International Cognitive Technology Conference (CT97)*.
- BUI, H. 2003. A general model for online probabilistic plan recognition. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- BURRIDGE, R. R., RIZZI, A. A., AND KODITSCHKEK, D. E. 1999. Sequential composition of dynamically dexterous robot behaviors. *Int. J. Robot. Resear.* 18, 6, 534–555.
- CHEN, Y., HUANG, H., XU, W., WALLIS, R., SUNDARAM, H., RIKAKIS, T., HE, J., INGALLS, T., AND OLSON, L. 2006. The design of a real-time, multimodal biofeedback system for stroke patient rehabilitation. *SIG ACM Multimedia*, Santa Barbara, CA.
- CHEN, Y., XU, W., SUNDARAM, H., RIKAKIS, T., AND LIU, S.-M. 2007. Media adaptation framework in biofeedback system for stroke patient rehabilitation. *SIG ACM Multimedia*, Augsburg, Germany.
- DARWICHE, A. 2001. Constant space reasoning in Dynamic Bayesian Networks. *Intl. J. of Approximate Reasoning* 26, 161–178.
- DEMPSTER, A., LAIRD, N., AND RUBIN, D. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc.* 39, 1, 1–38.
- DOURISH, P. 2001. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press Cambridge, MA.
- DUDA, R. O., HART, P. E., AND STORK, D. G. 2000. *Pattern Classification*. Wiley.
- FOX, M., GHALLAB, M., INFANTES, G., AND LONG, D. 2006. Robot introspection through learned hidden markov models. 170, 2, 59–113.
- GALLICCHIO, J. AND KLUDING, P. 2004. Virtual reality in stroke rehabilitation: Review of the emerging research. *Phys. Therapy Rev.* 9, 4, 207–212.
- GROUT, D. J. AND PALISCA, C. V. 2001. *A History of Western Music*. Norton, New York.
- GRUPEN, R. A. AND COELHO, J. J. A. 2002. Acquiring state from control dynamics to learn grasping policies for robot hands. *Advanced Robotics* 16, 5, 427–443.
- HE, X., MA, W.-Y., KING, O., LI, M., AND ZHANG, H. 2003. Learning and inferring a semantic space from user’s relevance feedback for image retrieval. *IEEE Trans. Circ. Syst. Video Tech.*
- HOLDEN, M. AND DYAR, T. 2002. Virtual environment training: A new tool for neurorehabilitation. *Neurology Rep.* 26, 62–72.
- HOLDEN, M., TODOROV, E., CALLAHAN, J., AND BIZZI, E. 1999. Virtual environment training improves motor performance in two patients with stroke: Case report. *Neurology Rep.* 23, 57–67.
- HOLLAN, J., HUTCHINS, E., KIRSH, D., AND SUTCLIFFE, A. 2000. Distributed cognition: Toward a new foundation for human-computer interaction research on the effective use and reuse of HCI knowledge. *ACM Trans. Comput.-Hum. Interact.* ACM Press. 7, 174–196.
- HSIAO, K., KAEHLING, L. P., AND LOZANO-PEREZ, T. 2007. Grasping POMDPs. In *Proceedings of the IEEE Conference on Robotics and Automation*.

- HUBER, P. J. 1981. *Robust Statistics*. Wiley.
- HUTCHINS, E. 1995. *Cognition in the Wild*. MIT Press, Cambridge, MA.
- ISHII, H. AND ULLMER, B. 1997. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press. 234–241.
- ISHII, H., WISNESKI, C., BRAVE, S., DAHLEY, A., GORBET, M., ULLMER, B., AND YARIN, P. 1998. ambientROOM: Integrating ambient media with architectural space. *CHI 98 Conference Summary on Human Factors in Computing Systems*. ACM Press. 173–174.
- JACK, D., BOIAN, R., MERIANS, A. S., TREMAINE, M., BURDEA, G. C., ADAMOVICH, S. V., RECCE, M., AND POIZNER, H. 2001. Virtual reality-enhanced stroke rehabilitation. *IEEE Trans. Neural Syst. Rehabi. Eng.* 9, 308–318.
- KAELBLING, L. P., LITTMAN, M. L., AND CASSANDRA, A. R. 1998. Planning and acting in partially observable stochastic domains. *101*, 99–134.
- KIRSH, D. 1995. The intelligent use of space. *73*, 31–68.
- KJAERULFF, U. 1995. dHugin: A computational system for dynamic time-sliced Bayesian networks. *Intl. J. Forecasting* 11, 89–111.
- LATOMBE, J. C. 1991. *Robot Motion Planning*. Kluwer Academic Publishers, Norwell, MA.
- LAVALLE, S. M. 2006. *Planning Algorithms*. Cambridge University Press.
- LIAO, L., FOX, D., AND KAUTZ, H. 2004. Learning and inferring transportation routines. In *Proceedings of the National Conference on Artificial Intelligence*.
- MA, Y.-F. AND ZHANG, H.-J. 2003. Contrast-based image attention analysis by using fuzzy growing. In *Proceedings of the 11th ACM International Conference on Multimedia*. ACM Press. 374–381.
- MAZALEK, A., DAVENPORT, G., AND ISHII, H. 2002. Tangible viewpoints: A physical approach to multimedia stories. In *Proceedings of the 10th ACM International Conference on Multimedia*. ACM Press. 153–160.
- MERIANS, A. S., JACK, D., BOIAN, R., TREMAINE, M., BURDEA, G. C., ADAMOVICH, S. V., RECCE, M., AND POIZNER, H. 2002. Virtual reality-augmented rehabilitation for patients following stroke. *Phys. Ther.* 82, 9, 898–915.
- MURPHY, K. 2002. Dynamic Bayesian networks: Representation, inference and learning. Ph.D. thesis. Dept. Computer Science, UC Berkeley.
- MURPHY, K., TORRALBA, A., AND FREEMAN, W. 2003. Using the forest to see the trees: A graphical model relating features, objects and scenes. In *Proceedings of the NIPS'03 (Neural Info. Processing Systems)*.
- PINEAU, J., GORDON, G., AND THRUN, S. 2003. Point-based value iteration: An anytime algorithm for pomdps. *International Joint Conference on Artificial Intelligence*.
- RABINER, L. R. 1989. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE* 77, 2, 257–286.
- RUSSELL, S. AND NORVIG, P. 2003. *Artificial Intelligence, A Modern Approach*. Pearson Education, Inc.
- SIMMONS, R. AND KOENIG, S. 1995. Probabilistic robot navigation in partially observable environments. *International Joint Conference on Artificial Intelligence*.
- SMITH, T. AND SIMMONS, R. 2004. Heuristic search value iteration for pomdps. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*.
- SPAAN, M. T. J. AND VLASSIS, N. 2005. Perseus: Randomized point-based value iteration for pomdps. *J. Artif. Intell. Resear.* 24, 195–220.
- SUN, Y., ZHANG, H., ZHANG, L., AND LI, M. 2002. A system for home photo management and processing. In *Proceedings of the 10th ACM International Conference on Multimedia*. 81–82.
- SUNDARAM, H. AND CHANG, S.-F. 2000. Determining computable scenes in films and their structures using audio-visual memory models. In *Proceedings of the ACM International Conference on Multimedia 2000*. 95–104.
- SUNDARAM, H. AND RIKAKIS, T. 2006. Experiential Media Systems. *Encyclopedia of Multimedia*. B. Furtht. Springer Verlag, NY.
- THEOCHAROUS, G. AND KAELBLING, L. P. 2004. Approximate planning in POMDPs with macro-actions. *Advances in Neural Information Processing Systems 16*, Vancouver.
- THEOCHAROUS, G., MURPHY, K., AND KAELBLING, L. P. 2004. Representing hierarchical POMDPs as DBNs for multi-scale robot localization. In *Proceedings of the International Conference on Robotics and Automation*.
- ULLMER, B. AND ISHII, H. 2000. Emerging frameworks for tangible user interfaces. *IBM Syst. J.* 39, 915–931.
- WEISER, M. 1993. Some computer science issues in ubiquitous computing. *Comm. ACM* 36, 75–84.
- WHITE, D., BURDICK, K., FULK, G., SEARLEMAN, J. AND CARROLL, J. 2005. A virtual reality application for stroke patient rehabilitation. In *Proceedings of the IEEE International Conference on Mechatronics & Automation Niagara Falls, Canada*.

- XIE, L., CHANG, S.-F., DIVAKARAN, A., AND SUN, H. 2003. Unsupervised discovery of multilevel statistical video structures using hierarchical hidden Markov models. In *Proceedings of the IEEE Conference on Multimedia and Expo 2003*. Baltimore MD.
- XU, W., CHEN, Y., SUNDARAM, H., AND RIKAKIS, T. 2006. Multimodal archiving, real-time collaborative annotation and information visualization in a biofeedback system for stroke patient rehabilitation. In *Proceedings of the 3rd Workshop on Capture Archival, Retrieval of Personal Experiences, in Conjunction with ACM MM 06*, Santa Barbara, CA.

Received February 2008; revised June 2008; accepted June 2008