An Adaptive Dialogue System for Assessing Post Traumatic Stress Disorder

Alexandros Papangelis
UTA Research Institute &
National Center for Scientific
Research "Demokritos", IIT
alexandros.papangelis@mavs.uta.edu

Robert Gatchel
Department of Psychology,
University of Texas at Arlington
gatchel@uta.edu

Vangelis Metsis
Department of Computer
Science and Engineering,
University of Texas at Arlington
vmetsis@gmail.com

Fillia Makedon
Department of Computer
Science and Engineering,
University of Texas at Arlington
makedon@uta.edu

ABSTRACT

In this paper, we present a system which is able to interact through natural dialogue, with PTSD patients, as well as to guide the conversation aiming to elicit enough information to make an assessment of their condition, in a manner similar to a self assessment test. Our system is able to adapt to each individual patient and can operate in two modes: one that stores information about previous sessions with a patient to provide a sense of trust and relationship; and one that does not store information to preserve anonymity.

Categories and Subject Descriptors

J.3 [**Computer Applications**]: Life and Medical Sciences; H.1.2 [**User/Machine Systems**]: Software Psychology

General Terms

Human Factors, Measurement

Keywords

PTSD assessment, adaptive dialogue, complex actions

1. INTRODUCTION

Post-Traumatic Stress Disorder (PTSD), according to the National Institute of Mental Health [7], refers to experiencing an anxiety disorder due to some traumatic event. It affects about 5% of the US population (or approximately 13 million people), according to the Sidran institute, with approximately 70% of Americans having experienced PTSD at least once in their lifetime. Moreover, women are twice as likely as men to develop PTSD and the economic burden incurred by treating PTSD is estimated to be $42.3bn annually. This disorder can occur at any age and under many circumstances, such as the person or a loved one experiencing sexual or physical assault, war or disaster, and other stressful events. The main symptoms of PTSD are reliving the event, avoiding anything that reminds one of it, and exhibiting hyperarousal. It is typically treated by cognitive-behavioral therapy (CBT) or medications. CBT involves sessions with a health professional which may include exposure therapy, cognitive restructuring or stress inoculation therapy. Medications for PTSD typically include antidepressants, antipsychotics and medications for helping patients relax and sleep [7]. Besides visiting a doctor, one can take self-assessment tests, either online or in an appropriate facility, that may yield a measure of PTSD symptomatoly. Online methods have also been proposed and are currently available, such as the SimCoach [11], which is an online virtual agent aiming to motivate military personnel to seek healthcare. Other approaches include online self-assessment tests[1], and chat rooms [2] or forums [3] where the interested person may find information and support. A critical factor not yet addressed in-depth by the above automatic systems is that the patient needs to trust the system in order to confide his / her experiences, and possibly accept treatment. The patient can, however, be assumed motivated to tell the truth, because he/she has reached out for help. In order to achieve trust between a patient and an automated system, therefore, we need to create a patient profile where information about preferences, previous sessions, diagnoses, etc. may be stored, such as [11]. This means that the patient must be assigned a patient name and password and thus log into the system using some form of id. Because this may be a hindering factor to many patients, especially if they are not familiar with the system, we propose two modes of operation: one that requires logging in, where the system can remember information regarding previous sessions; and one that does not require logging in and the system does not store any information. Realistic avatars also help cre-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
PETRAE May 29 - 31, 2013, Rhodes, Greece
Copyright 2013 ACM 978-1-4503-1300-1/12/06 ...$15.00.

[1]www.online-therapy.com/ptsd-test-online-c-155_161.html
www.healthyplace.com/psychological-tests/ptsd-test/
[2]www.healthfulchat.org/ptsd-chat-room.html
[3]www.ptsdforum.org/c/portal/

ate trust and a sense of personal relationship. The novelty of our approach is that we keep track of the patient's emotional state and generate content in real-time in an attempt to avoid unpleasant behaviors and adapt to each individual patient. It should be noted that there is no more risk in using our system than the risk when taking an online self-assessment test. We do not propose to practice therapy, but to only assess the patients' condition, using well-defined and accepted methods. In the following Section we present background information; in Section 3 we describe our system; in Section 4 we present our PTSD assessment method; and we conclude in Section 5.

2. BACKGROUND KNOWLEDGE

In order to represent the interaction and be able to "solve" the dialogue problem (i.e. to elicit the necessary information from the patient), we follow the *Information Seeking* paradigm [10]. This paradigm is based on the assumption that the system needs some pieces of information in order to perform the tasks it has been designed to perform. In our case, for example, in order to assess the patient's condition, the system needs some pieces of information that can be provided through question and answering. Each piece of information therefore is called a *slot* and, in our modeling, can take discrete values, such as 'yes', 'no', 'once', etc. More specifically, we apply the model described in [8], according to which we have a set of N slots $Z =< z_0, ..., z_N >\in V, V = V_0 \times V_1 \times ... \times V_N, V_i = \{1, ..., T_i\}$, where $z_i \in V_i$. We also have a non-empty set of system actions A that, in our system, are defined as prompts for information. System actions can be either basic or complex (i.e. combinations of other actions). We will provide more details about complex actions and complex-action learning in the following Section. Apart from system actions, we define U as the set of possible patient actions, meaning the type of information the patient may convey (e.g., a 'yes' or 'no' answer or a description of an event). The state of the interaction is captured by the dialogue state, which is defined as a vector $d \in D$, which contains all the necessary information to describe the interaction so far, such as pieces of information provided and pieces of information still missing, etc. As the dialogue progresses and the system gathers information, some actions may not make sense to be available, or may not be feasible. To take this into account, the model defines an availability matrix $\tilde{A} \in \{0, 1\}^{|D|\times|A|}$ where an action a is available at (dialogue) state d if $\tilde{A}(d, a) = 1$. Last, we define dialogue state transition probabilities, which are continually updated, in order to account for uncertainty in understanding the patient's utterance and we assign confidence values for each slot value (piece of information retrieved from the patient). To solve this problem, we model it as a hierarchical Markov Decision Process, and apply online hierarchical Reinforcement Learning techniques.

3. ADAPTIVE DIALOGUE SYSTEM (ADS)

In this Section, we present our system, targeted for interacting with PTSD patients and making an assessment of their condition. In this first version of the system, the patient must provide his / her response in spoken or textual form, and the system replies using both text and speech. In the future, we plan to add visual input and extract more audiovisual features, such as pitch or tone and facial expressions which lead to better emotion recognition. This will also enable us to detect if pauses or delays are due to external events (e.g., the patient's attention is diverted by someone), or other factors related to his / her condition. Currently, our ADS is able to guide the conversation in such a way that it elicits information similar to the information contained in a PTSD self-assessment test. The system continually monitors the patient's emotional state, and balances between keeping the patient calm or happy (if possible) and retrieving the information it needs by asking the appropriate questions. In order to achieve adaptation to each individual patient, we employ online complex-action learning and goal achievement techniques, described in [9]. These techniques allow the system to learn how to achieve complicated tasks, such as "analyseSumptom" by combining simple actions, such as "askFrequency" and "askIntensity". This technique also allows the system to adapt to each patient by altering the way complicated tasks are achieved (i.e., by asking different questions or rephrasing them). Goal-achievement techniques allow the system to "stay on target" which, in our case, means eliciting information necessary for PTSD assessment and keeping the patient at least in a calm state. If the patient's emotional state "worsens," the system will attempt to correct it by providing encouragement or, in future versions, talking about something different (e.g., sports news before coming back to PTSD). We make the assumption that the patient's input is categorized into answers to system questions, asking questions, or describing an event (which may or may not be relevant to the conversation). We provide more details on event recognition in the next Section. Out of these three categories, we can define the possible patient actions as U =<*ANSWER*, *ASK*, *DESCRIBE*, *END*>. It should be noted that an *ANSWER* action refers to answers to very specific questions, such as "What is your age?" or "How frequently does this happen?", while a *DESCRIBE* action refers to answers to more open-ended questions, such as "What is troubling you?" or "What happened that day?". *END* means that the patient intends to end the conversation. The actions available to the system are: A =<*GREET*, *INTRODUCE*, *ASK-YES-NO*, *ASK-OPEN*, *GIVE-FEEDBACK*, *CONFIRM*, *REPEAT END*>. *GREET* and *INTRODUCE* means that the system greets the patient and introduces itself, respectively. *ASK-YES-NO* means that the system asks a question that needs a 'yes' or 'no' answer, according to our self-assessment test model [1]. *ASK-OPEN* refers to asking an open-ended question, requesting elaboration on an event. *GIVE-FEEDBACK* provides encouragement or feedback to the patient about his / her condition. *CONFIRM*, *REPEAT* and *END* refer to asking the patient to confirm something s/he mentioned, asking the patient to repeat the last utterance and ending the interaction. Using this configuration, we are able to elicit, for example, frequency of occurrence information by either asking many *ASK-YES-NO* questions, such as "Does this occur more than once a week?" or one *ASK-OPEN* questions, such as "How often does this happen?", according to the state of the interaction, the current patient, etc.

3.1 Complex Action Learning

System actions can be further categorized into *basic* or *complex*. Basic actions refer to the system actions mentioned earlier, while complex actions refer to combinations of those actions into hierarchical structures that are able to

Complex Actions	Basic Actions
ASK-YES-NO-CONF	ASK-YES-NO CONFIRM
RETRIEVE-EVENT	GREET ASK-OPEN
ANALYSE-EVENT	RETRIEVE-EVENT CONFIRM
IDENTIFY-SYMPTOM	ASK-YES-NO RETRIEVE-EVENT ANALYSE-EVENT
RETRIEVE-SYMPTOM	RETRIEVE-EVENT IDENTIFY-SYMPTOM

Table 1: Example Complex Actions.

solve more complicated tasks, such as "retrieveSymptoms". Many complex tasks have common subtasks, so complex-action learning allows us to re-use knowledge of how to achieve a task in other more complicated tasks. Table 1 presents some examples of complex actions that our system can learn, by combining basic or other complex actions. In order to learn how to combine basic actions into complex ones, we apply the Action Weights Learning (AWL) method, described in [9], which is able to rank basic and complex actions according to their performance, and create an optimal set of complex actions. This set depends on the feedback the system receives from the environment and the patient and, therefore, the complex-action set for each patient may vary, meaning that, for each patient, the way to achieve a complex task may vary. Applying online complex-action learning has many advantages, such as the fact that we are able to re-use knowledge of how to solve a subtask, which saves time when encountering new complicated tasks. Moreover, complex-action learning allows for better adaptation to each individual patient's needs, and helps make the system appear more intelligent which can be critical in forming long-term relationships with the patients.

3.2 Goal Achievement

Our system is designed to elicit information from a PTSD patient and provide a basic form of assessment of the patient's condition. To achieve this, we have identified two goals that must be achieved: (a) the system needs to elicit enough information to make the assessment; and (b) the system needs to keep the patient from getting too frustrated, sad, angry, etc. The second goal must be met at all times (i.e. the system must monitor the emotional state of the patient and make sure the patient is at least calm). The first goal is also important because not meeting it would defeat the purpose of the system. These are classified as mandatory goals. We can also have optional goals, such as keeping the patient happy rather than calm, or retrieving extra information. Of course, meeting such optional goals may not be always possible, but they will be met when possible (perhaps by changing the order in which questions are asked that may affect the emotional state of the patient). In order to make sure the system always achieves its goals, we employ a technique called User State Estimation, also proposed in [9]. This technique provides enough flexibility to define mandatory and optional goals, and it guarantees that mandatory goals will be met (as long as a way to achieve them exists), and optional goals will be met when possible.

3.3 Natural Language Processing

In order to understand the patient's spoken input, we used the Sphinx Language Model Tool [12] to generate a dictionary and language model, based on a vocabulary we created composed of words relevant to our system. We then used pocket Sphinx [6] for Automatic Speech Recognition, for patients who choose to use a microphone. We opted not to use a complete vocabulary of English in order to make understanding easier, and to directly cut-off input that does not contain useful information. For example, if the patient is describing an event that occurred in the past, we are only interested in the fact that it is an important event (which may lead to the identification of a PTSD symptom), and on some keywords hinting of its severity. In future versions of our system, we plan to apply a comprehensive NLU component that may uncover more information about the patient's condition. In order to generate the system's output in natural-language, we used simpleNLG [5], which is a publicly available natural language generator. In order to use it, we need to provide parts of speech, such as verb, subject, object, etc. In the current version of the system, each system action has a predefined set of parts of speech, which generate the system output. In the future, however we will take many things into account, such as adaptive NLG and personalization. We, however, adapt the amount of encouragement we provide to the patient, depending on the estimate of his/her emotional state. We currently achieve this in a very simple way, we have five levels of encouragement and, as the patient's emotional state "worsens," the system gradually increases the amount of encouragement provided. For example, the system, instead of just asking a yes/no question, provides some statistical facts, such as "13 million other Americans currently have PTSD". We also mark the patient's response to different types (e.g., providing statistics, explaining details about PTSD, etc) and levels of encouragement, and then store it for future sessions.

3.4 Emotion recognition

In order to recognize and keep track of the patient's emotional state, we applied the textual emotion-recognition method proposed by [3], using emotional keywords identified by [4]. Because the vast majority of keywords in this vocabulary refer to single emotional states, this results in a relatively simple way to recognize the patient's emotional state. In the future, however, we intend to apply more complex methods that include multiple modalities (text, speech, video). In order to represent the emotional state, we follow the example of [3] and use a vector $e \in \mathbb{R}^{10}$ where each dimension corresponds to the emotional states identified by [4]: Happiness, Caring, Depression, Inadequateness, Fear, Confusion, Hurt, Anger, Loneliness, Remorse. The value at each dimension corresponds to the intensity of that emotion: $e_i \in [0, 1], i = 1..10$. We update the patient's emotional state estimate immediately after receiving input, taking into account the current recognized emotional state, as well as the previous estimate. In short, according to the method proposed by [3], we look for emotional keywords in the input sentence, and average the intensities of each emotion to come up with the current emotional state. This emotional state is multiplied by a weight, as a consequence of emotion *modification words*, such as 'very', 'extremely' or 'not'. At each point in time, therefore, we have a vector of emotional state intensities that correspond to the overall emotional state.

4. PTSD ASSESSMENT

To perform basic PTSD diagnosis, we collect information similar to the information a self-assessment test would collect, administering the PTSD Checklist [13] to calculate a score. Based on this score, we subsequently make an assessment of the patient's condition and provide appropriate feedback. In the rest of this Section, we describe: how we attempt to affect the patient's emotional state by providing appropriate encouragement; how we identify critical events and process them to see if they resemble PTSD symptoms; and, if so, how we assess their intensity and frequency of occurrence. We currently use a very simple method for identifying traumatic events, and that is by making the very strong assumption that the patient will describe or talk about an event when the system inquires about it. We then look for specific keywords in the patient's response to identify possible PTSD symptoms. In order to identify frequency of occurrence, we again assume that the patient willingly answers the system's questions. To get an estimate of the intensity of an event, we use the emotion recognition module and calculate a normalized score $intensity \in 0, ...4$, to correspond to the CAPS. According to [13], an event must be of moderate intensity or more, and occurs at least once or twice a month to be considered as a symptom. To make an assessment regarding PTSD, we use the rule described in [13] which states that, in order to diagnose PTSD, we need: one symptom of reliving a stressful event; three symptoms of avoiding situations that remind the patient of the event; and two symptoms of hyper-arousal. If these criteria are not met, the patient is classified as no-PTSD. Otherwise, we classify him / her into 3 categories, namely mild-PTSD, moderate-PTSD or severe-PTSD, according to a score that we calculate as follows: $score = \frac{i_t \times f_t}{\#t}$, where i_t is the intensity of symptom t and f_t the corresponding frequency of occurrence. At the end of the interaction, we provide feedback ("verdict") for each category, depending on the estimated severity of the condition. Also, during the interaction we provide encouragement, by providing statistical results, such as: the percentage of Americans that currently have PTSD or have had in some point in their lives; how long on average it takes to treat; and details on what PTSD is about ("fight or flight" reaction, etc.). We plan to develop a method for adapting the amount of encouragement that the system provides during the session and the final "verdict" the system provides at the end of the session, depending on the patient's profile (that is built throughout multiple sessions) if available, and his / her emotional state. Note that in the case of no-PTSD, we provide feedback giving general information about psychological disorders / encouragement, as we feel that if the patient went out of his / her way to interact with our system there may be either another condition we are not able to identify, or that we were unable to correctly diagnose PTSD.

5. DISCUSSION AND FUTURE WORK

In the future, we plan to incorporate visual input for the patients that choose the version that requires a log in and also are willing to provide such information. This will enable us to detect facial expressions and, combined with acoustic features such as pitch and tone of the patient's voice, we can have better emotion recognition and lower false positives rates as, for example, a long delay to respond may be due to another person diverting the patient's attention. We also plan to apply tools such as the Natural Language Tool Kit [2] that will provide us the semantics of the patient's utterance, aiming to achieve a more natural interaction with the patient and possibly talking about something different than PTSD (e.g., sports) for some time in order to form a long-term relationship with the patient. This system will be part of a larger companion that will operate in assistive-living environments and, while its purpose will not be assessment or diagnosis, it will keep an eye for behavioral cues, indicating possibly serious conditions, such as depression or PTSD.

6. ACKNOWLEDGEMENTS

This work is supported in part by the NSF under award numbers CNS 0923494, CNS 1035913 and IIS 1238660.

7. REFERENCES

[1] Diagnostic and statistical manual of mental disorders, 1994. 4th Edition. American Psychiatric Association.

[2] S. Bird, E. Klein, and E. Loper. *Natural language processing with Python.* O'Reilly Media, Incorporated, 2009.

[3] W. C. Chuang, Z.J. Multi-modal emotion recognition from speech and text. *Computational Linguistics and Chinese Language Processing*, 9(2):45–623, 2004.

[4] T. Drummond. Vocabulary of emotions.

[5] A. Gatt and E. Reiter. Simplenlg: A realisation engine for practical applications. In *Proceedings of the 12th European Workshop on Natural Language Generation*, pages 90–93. Association for Computational Linguistics, 2009.

[6] D. Huggins-Daines, M. Kumar, A. Chan, A. W. Black, M. Ravishankar, and A. I. Rudnicky. Pocketsphinx: A free, real-time continuous speech recognition system for hand-held devices. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 1, pages I–I. IEEE, 2006.

[7] U. S. D. of Health & Human Services. Post-traunatic stress disorder. *National Institute of Mental Health*, (08-6388), 2008.

[8] A. Papangelis. A comparative study of reinforcement learning techniques on dialogue management. In *Proc. of SRW at EACL*, pages 22–31. ACL, 2012.

[9] A. Papangelis, V. Karkaletsis, and F. Makedon. Online complex action learning and user state estimation for adaptive dialogue systems. In *Proc. of IEEE ICTAI*, 2012.

[10] O. L. V. Rieser and O. LEMON. Reinforcement learning for adaptive dialogue systems. 2011.

[11] A. Rizzo, K. Sagae, E. Forbell, J. Kim, B. Lange, J. Buckwalter, J. Williams, T. Parsons, P. Kenny, D. R. Traum, J. Difede, and B. Rothbaum. Simcoach: An intelligent virtual human system for providing healthcare information and support. In *I/ITSEC 2011*, Orlando, Florida, November 2011 2011.

[12] A. Rudnicky. Sphinx knowledge base tool (2010).

[13] F. W. Weathers, T. M. Keane, and J. R. Davidson. Clinician-administered ptsd scale: A review of the first ten years of research. *Depression and anxiety*, 13(3):132–156, 2001.