Emotion Expression in a Socially Assistive Robot for Persons with Parkinson's Disease

Andrew Valenti andrew.valenti@tufts.edu Tufts University Human-Robot Interaction Laboratory Medford, MA Avram Block Tufts University avram.block@tufts.edu Meia Chita-Tegmark Tufts University mihaela.chita_tegmark@tufts.edu

Michael Gold Tufts University michael.gold@tufts.edu

ABSTRACT

Emotions are crucial for human social interactions and thus people communicate emotions through a variety of modalities: kinesthetic (through facial expressions, body posture and gestures), auditory (the acoustic features of speech) and semantic (the content of what they say). Sometimes however, communication channels for certain modalities can be unavailable (e.g., in the case of texting), and sometimes they can be compromised, due to a disorder such as Parkinson's disease (PD) that may affect facial, gestural and speech expressions of emotions. To address this, we developed a prototype for an emoting robot that can detect emotions in one modality, specifically in the content of speech, and then express them in another modality, specifically through gestures.

The system consists of two components: detection and expression of emotions. In this paper we present the development of the expression component of the emoting system. We focus on its dynamical properties that use a spring model for smooth transitions between emotion expressions over time. This novel method compensates for varying utterance frequency and prediction errors coming from the emotion recognition component. We also describe the input the dynamical expression component receives from the emotion detection component, the development and validation of the output comprising of the gestures instantiated in the robot, and the implementation of the system. We present results from a human validation study that shows people perceive the robot gestures, generated by the system, as expressing the emotions in the speech content. Also, we show that people's perceptions of the accuracy of emotion expression is significantly higher for a mass-spring dynamical system than a system without a mass-spring when specific detection errors are present. We discuss and suggest future developments of the system and further validation experiments.

This paper is part of a larger project to develop a prototype for a socially assistive robot for PD persons. The goal is to present the

PETRA '20, June 30-July 3, 2020, Corfu, Greece

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7773-7/20/06...\$15.00

https://doi.org/10.1145/3389189.3389190

Matthias Scheutz Tufts University matthias.scheutz@tufts.edu

technical implementation of one robot capability: emotion expression.

CCS CONCEPTS

• Computing methodologies → Information extraction; Classification and regression trees; • Computer systems organization → Robotic components; • Human-centered computing → Collaborative and social computing.

KEYWORDS

machine learning for robotics; sentiment analysis; topic modeling; socially interactive agents; human-robot interaction and collaboration; dynamical systems

ACM Reference Format:

Andrew Valenti, Avram Block, Meia Chita-Tegmark, Michael Gold, and Matthias Scheutz. 2020. Emotion Expression in a Socially Assistive Robot for Persons with Parkinson's Disease. In *The 13th PErvasive Technologies Related to Assistive Environments Conference (PETRA '20), June 30-July 3, 2020, Corfu, Greece.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3389189.3389190

1 INTRODUCTION

People communicate emotion using multiple modalities, such as tone of voice, facial expressions and gestures. However, in some situations, one or more modalities may be absent, noisy, or damaged and this may degrade how the human body expresses emotions. Because people rely heavily on facial expression in attributing and interpreting other's emotions and motivational states, compromised or missing modalities can deeply affect the person's ability to communicate which may lead to impaired social interactions and reduced quality of life. Such is the case for people living with Parkinson's disease (PD) who, due to a condition called facial masking, are impaired in their ability to express their inner emotional state.

Our long-term goal is to develop a robot that could help people with PD express their inner emotional state and thus improve their communication with caregivers. Our overall research objective, which is the next step in attainment of our long term goal, is to develop an emoting system for robots that is able to detect an emotion in an unaffected modality, namely content of speech and express it in an absent one, namely gestures. We suggest that an emoting robot could help people with PD express their emotional states and thus improve their communication with caregivers. In

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

	Overall project	This paper
Objective	Develop a socially assistive robot for the social self- management of health of people with PD	Develop a prototype for an emoting robot that can de- tect emotions in one modality (content of speech) and
	management of neuron of people while 2	express them in another (gestures).
Robot capabilities	Simple interaction, observation of activities, mediation of interactions.	Detection and expression of emotion.
Hypothesis	The assistive robot will reduce stigma and improve com- munication between people with PD and caregivers and health-care providers.	Robust emotion expressions of the robot can be cor- rectly perceived in both high and low frequency emot- ing conditions.
Validation	Clinical trial with people with PD and their caregivers.	General population studies with robot emoting in dif- ferent conditions.

Table 1: Project scope.

this paper, part of a larger project (see Table 1), we describe the emoting system which uses two components: an emotion detection component and an expression component. This paper focuses on the latter. We extended the unsupervised emotion prediction model described in [38] and trained it to detect five different states of emotional valence. The detected emotional state drives a massspring dynamical system to smooth emissions from the detection component, e.g., compensating for varying utterance frequency and prediction errors. Detected emotions are expressed as gestures in the Nao robot from SoftBank Robotics. The Nao robot has been previously used in therapy for other groups of people, such as individuals with autism [4, 10, 17, 31]. The mass-spring also ensures that relatively more emotional "force" is needed to move a person from a more extreme emotional state than from a more neutral state. To test our system we conducted three experiments with human participants.

This paper proceeds as follows. In section 2, we talk about the prior approaches used to detect emotional content in text, review prior work in applying dynamical systems to emotional modeling, and discuss how emotions are embodied in humans and robots. Section 3 discusses the development and implementation of the mass-spring dynamical system. In section 4, we explain the three human-robot interaction experiments we ran to validate the full system and to highlight the contribution of the mass-spring. The results of the study show that when embedded in the Nao robot, participants recognized the connection between the robot's gestures and the speaker's utterances more when the robot emoted based on the model's predictions than when the robot emoted randomly. Additionally, the mass-spring dynamical system led to greater perceived association between the robot's gestures and the emotional content of the speaker's utterances than a model without a massspring element, when the emoting was done at a low-frequency (for every third utterance). This indicates that the mass-spring dynamical system is more robust to errors. Finally, in sections 6 and 7, we discuss the advantages, disadvantages, and limitations of this approach and further improvements that can be made to the system.

2 BACKGROUND

2.1 Emotional Modeling

Sentiment analysis most often refers to the techniques used to infer the binary emotional polarity (e.g., positive, negative) of a person as they interact with a text or a document, rather than determining the specific human emotion [20]. However, sentiment analysis approaches may also attempt to classify additional emotion labels. Rather than recognizing emotions as belonging to discrete and often binary categories, *emotion recognition* attempts to infer a set of emotion labels such as happiness or satisfaction, both of which fall under the positive category. The following overview covers detecting emotions in text; emotion modeling is much wider and includes using speech, gestures, facial expression modalities either singularly or in combination to detect emotions.

The emotion labels to be detected computationally in text may be derived from the various psychological theories of emotion. For example, Eckman [8] argues that humans share six basic emotions: happiness, sadness, fear, anger disgust, and surprise. Plutchick [24] identified eight primary emotions: anger, anticipation, joy, trust, fear, surprise, sadness and disgust, some being opposites (e.g. joysadness) and some lending themselves to combinations and various intensities. This trend towards measurability and continuity in emotional models was solidified by Russell's Circumplex Model of Affect[27]. He arranges emotions in a circle around two axes: arousal and valence. This creates a continuous two dimensional space in which emotions can be plotted. Thus, some computational models attempt to detect continuous values of emotional valence or arousal e.g., [16, 33, 35].

Existing approaches to emotion recognition and sentiment analysis fall broadly into three categories: statistical, sentiment lexicon, and a combination of the two [3]. The challenge for the statistical approach is to find document features which will be sufficiently discriminating so that a classifier can separate the document into the desired sentiment categories. Deriving these features from the speech signal directly as in [29, 30] has the advantage of using both affect in the voice and in the content; however, this assumes an intact vocal channel which is not the case in persons with PD. Regardless of the modality, the challenge is to discover a set of discriminative features and it is often difficult to discover them a priori. As a result, features are often found through experimentation using various selection methods such as information gain [19], principal component analysis for speech, and manually selected emotion keywords for text [5]. In these cases, feature selection is tied to a particular database which may result in poor generalization. A principled way for feature selection would overcome these limitations by utilizing features that operate at a higher level of abstraction, one that is common across text or speech datasets. For example, Shah [30] used Latent Topic Models to extract emotionally-salient features from speech; VADER [16] and LIWC [23, 35] use a sentiment lexicon and rule-based approach to detect emotion in text. More recently, results exceeding sentiment analysis benchmarks have been reported by "fine-tuning" a pre-trained language model such as the Bidirectional Encoder Decoder Transformer (BERT) for aspect-based sentiment analysis [34]. To detect and characterize the emotional content of communications from individuals with PD, the prevalent approach employed so far has been to use a sentiment lexicon such as LIWC [35]. However, Valenti et al. [40] suggested that LIWC did not categorize positive and negative emotion contained in a document as well as a topic modeling approach when the word count approached that of the average sentence, i.e., 13 words.

2.2 Dynamical Systems of Emotion

Prior research suggests physical models, in particular the massspring, can be used to simulate physical human movement and create "plausible" behaviors [9]. Given their capacity to simulate physical behaviors, mass-spring models can also be used to simulate human movement qualities. Thus we chose to use the mass-spring model to simulate human gesture movement when transitioning from one emotion to another. Without this component, these movements would be completely discrete, and would not give rise to the continuous transitions between emotion states that we observe in humans. The goal of this component is to initiate behaviors in the Nao which imply transitions through intermediate emotional states on the way to more extreme goal states, or on the way back to neutrality from these extremes.

Previous studies on dynamic emotional models support the presence of continuity. Peng et al. [43] model emotional transitions according to the probabilities that arise in transitions between emotion states. The study shows that in a given emotional state, some subsequent states are far more likely than others. Specifically, under positive external influence, the researchers claim that negative emotions are most likely to transition to neutral emotions, and neutral emotions are most likely to transition to positive ones. In a discrete model based purely on the affect of human speech, these necessary intermediary steps are skipped and negative states may be forced to transition directly to positive ones. Similarly, an adaptive emotional model produced by Han et al. [13] was developed by mimicking the emotional behavior of a human agent over the course of a conversation spanning seven emotion-specific dialogues. When charted along the two-axis Circumplex model, it is clear that their agent undergoes continuous, incremental progression as it moves toward the appropriate emotional region during each of these dialogues.

In addition, the mass-spring model ensures continuity using the theory of emotional decay proposed by Velasquez in [42]. In this theory, an emotional state is not maintained for the exact duration of the triggering stimulus and dropped when the stimulus disappears. Rather, the onset of a stimulus instigates the development of an emotional state, which slowly reverts to neutrality over time. The application of this theory can also be observed in the model developed by Yang et al. via a different implementation than the one proposed here [44].

2.3 Emotional Embodiment in Humans and Robots

The challenge is to design the robot's behavior so that it expresses the detected emotion in a naturalistic way which compensates for the lack of facial and gestural affect cues in the persons with PD. Research has shown that humans can successfully estimate an agent's emotional state purely from their body movements [6]. For this study, we used the Nao robot to display the emotional state detected from continuous speech through body movements. A number of studies of human affective body language recognition as well as robotic affective body language production have helped ground our approach.

Studies regarding emotional body language fall into two categories: attempts to create models by which robots are able to estimate the emotional state of a human agent through visual tracking, and attempts to assess how the same neutral action is executed differently depending on the affect of the agent. In both cases, the emotions examined were often a specific subset of the emotions described by Russell's Circumplex model. We chose to represent only the valence axis of emotion, so the gestures designed are meant to represent gradations of positivity or negativity, rather than specific emotions. For this reason, the poses and movements observed and produced in these other studies could not be replicated directly, but do serve as the basis for our robot's behaviors. In their research, Shan et al. use the FABO video database to develop an algorithm to assess human emotion from body gestures. This research provides evidence for the recognizability of raised arms and hands as an indicator of joy and excitement [32]. Another study, which also used a Nao robot, confirms the recognizability of happiness through raised hand and arm gesticulation [21]. Similarly, research by de Silva and Bianchi-Berthouze analyze the salience of various body-feature point collections as a method of quantitatively describing body language to develop a classifier for emotionally labeled body language performed by an actor. This data confirms the recognizability of gestures such as drooping chest and raising hands towards one's face as an indicator of sadness [7]. These studies provided static poses which could be generalized, and converted into animated body movements, to emulate varying degrees of positivity.

Other studies showed how specific poses or motions can be altered to convey a given emotion. One study experimented with how the upward or downward angle of the robot's head can help clarify the emotion it displays. The results of this study suggested that a down-turned head and face helped convey fear and sadness, while an upturned head and face helped convey pride, happiness and excitement. We use this in our gestural design as an increase in head angle change in the appropriate direction to indicate increased negativity or positivity. Amaya et al. [2] developed an algorithm for determining the physiological effect of performing the same action (e.g., drinking from a cup or kicking a ball) while conveying neutral, sad or excited affects. They determined that when excited, an action is generally performed with higher joint velocity and more direct movements, whereas when sad, the joints involved move slower and less efficiently.



Figure 1: Within DIARC, the Emotion Pipeline consists of the prediction (Detector) and dynamical system (Expressor) components. It receives utterances from the Automated Speech Recognizer (LVASR) and sends its predicted affect to the Goal and Action Manager

3 DEVELOPMENT OF THE MASS-SPRING DYNAMICAL SYSTEM

3.1 System Input

In Figure 1, we show the major building blocks of the robot's emotional regulation and natural language understanding systems of the DIARC cognitive robotic architecture, of which a detailed description is given in [28]. The figure shows how we supplemented DIARC with an emotion pipeline consisting of two components, the *Detector* and the *Expressor*, along with their incoming and outgoing connections. The Detector component receives a text utterance from the Large Vocabulary Automatic Speech Recognizer (LVASR) in the Perception layer and gives its prediction to the Expressor, discussed in Section 3.3. The Expressor component then sends the goal predicates representing the desired affective states to the Goal Manager.

The LVASR is based on the chain model developed for the ASpIRE Challenge [14] and trained on Fisher English that has been augmented with impulse responses and noises to create multi-condition training [41]. The chain model uses Kaldi, a toolkit for speech recognition written in C++ and licensed under the Apache License v2.0 [25]. Kaldi has demonstrated low error rates in a variety of challenging acoustic environments using conversational speech [1, 15]. The challenge for the LVASR is to determine the endpoints of the utterances, in which words are connected together instead of being separated by speech codes such as pauses. Unknown boundary information about words, co-articulation, production of surrounding phonemes, and rate of speech affect performance [12] and the LVASR may generate textual representations that vary depending on the speaker, e.g., disfluencies, speech rate. As a result, the continuous speech will likely generate utterance transcriptions from which the detector may generate an incorrect prediction. Furthermore, persons with PD may have long pauses or other speech anomalies depending on the disease progression. The mass-spring tries to smooth these prediction errors and compensate for different emission frequencies from the LVASR and predictor.



Figure 2: Individual gestures on the robot representing emotional valence. Top row shows three levels of increasing positivity, starting with the least positive on the left. The image in the center is neutral. The bottom row shows three levels of positivity increasing from neutral. The *M* mean and *SD* standard deviations of participant ratings are presented under each gesture.

3.2 System Output

In order to express the results of the emotion detector and massspring system in the physical world, it was necessary to define behaviors in the Nao robot which would best reflect the emotional data produced. All physical movement design took place in the Nao's companion software, Choregraphe. This software contains a method entitled *Animation Mode* which allows for manual moving and recording of each of the Nao's joints at various time steps. While the gestures defined here are generally based on observations of human gesticulation, prior research has produced specific data which was used to improve their efficacy and accuracy.

To produce reliable and easily interpretable gestures in the robot, we incorporated the prior research described in Section 2.3 as follows. Gestures meant to convey positivity incorporated raised arms and hands. This is supported by [32] as well as [21]. Gestures meant to convey negativity include drooped chest and hands raised towards the robot's face. This is supported by research such as [7]. Rotation of the head in upward or downward directions was also used to further convey positivity or negativity respectively. This feature is supported by the work presented in [18]. Finally, the speed at which a gesture was performed was manipulated. In following with research presented in [2], more positive gestures were

PETRA '20, June 30-July 3, 2020, Corfu, Greece

carried out with higher joint velocity and swift, direct movements, while more negative gestures involved slower, less direct motion. Still images displaying a frame from each animation developed through this process can be seen in Figure 2.

Having designed potential gestures for the Nao to execute in order to reflect a specific internal emotional position along a valence scale, we sought to ensure that these gestures successfully conveyed their target information. To accomplish this, we carried out a validation study on these gestures in isolation from the rest of the architecture and pipeline. We recorded 10 second videos of each of the seven gestures show in Figure 2 which we showed to 25 human subjects (28% Female; Mean age = 32.08, SD = 9.39) who participated in the validation study on Amazon Mechanical Turk. The videos were presented in random order and four foil videos, of gestures from the standard Choregraphe database, were interspersed in between. After watching each video, participants answered the following question: "How positive or negative are the emotions expressed by the robot in this video?" The answer was given on a seven-point Likert scale from "Very Negative" to "Very Positive" with a middle point labeled "Neutral". We note that each video was rated independently, which is a more stringent test than simply ranking-ordering the videos from most negative to most positive. Means and standard deviations of ratings for each gesture are shown in Figure 2. When averaging across the positive gestures and across the negative gestures separately, the positive gestures are rated as expressing significantly more positive emotions than the negative ones, t(24) = 8.90, p < 0.001.

3.3 Mass-Spring Dynamical System



Valence Axis of Emotion

Figure 3: Mass-spring model with emotion particle at neutral valence. Applying a positive force F compresses the spring to move the particle toward a stronger positive valence; applying a negative force stretches the spring, moving the particle in a less positive direction. In the absence of a force, the spring's restoring effect will move the particle towards neutral.

We have attempted to replicate the physical behavior of human expression by modeling the Nao's emotional *position* along the valence axis of emotion as a particle. In the current study, this particle behaves as the mass in a mass-spring physical system (see Figure 3). This physical model was selected because it aligns well with the desired behavior of this particle. In the mass-spring system, an external force applied draws the mass away from its neutral resting position and the restorative force of the spring pulls it back. Additionally, the further that the particle is from its resting position, the stronger the external force must be to increase this distance. The model we used is represented by the following second order differential equation:

$$\ddot{x} = \frac{-kx - b\dot{x}}{m} + F \tag{1}$$

where

- \ddot{x} = the acceleration of the particle
- *k* = the spring constant, which defines how easily the spring is stretched or compressed
- *m* = the mass of the particle, which in the current research is set to 1
- *x* = the position of the particle, relative to its resting position; this is the value which is tracked internally to maintain current emotional state
- *F* = the force applied to the particle
- b = the damping constant applied to the velocity \dot{x}

The basic mass-spring model results in a system in which the particle undergoes infinite oscillation between resting position and the farthest distance attainable under a given external force. As this fluctuation does not reflect standard human emotional behavior, we incorporated a damping force into the model. The damping constant, *b*, minimizes this oscillation by causing the particle to lose energy as it returns to its resting point. Critical damping (b_c) , in which no oscillation occurs whatsoever, is attained when $b = 2\sqrt{km}$. This is the formula used to determine the value of the *b* term in Equation 1.

A number of the parameters provided to this model, such as k, m, and b, have been determined through in-lab experimentation and tuning in order to produce consistent and expected behaviors in the system. Future research which expands on the utility of applying physical system constraints to emotional or social behavior may include a more focused study on determining the optimal parameter set for the mass-spring system to produce natural behavioral transitions.

3.3.1 Mapping Categorical Predictions to Force. Equation 1 depends on the force F which acts on the mass to move it in the positive or negative direction, depending on the valence classification of speech. The Detector component generates five possible classifications, of increasing positivity. Calling get_prediction_value() method of TopicModel (see Figure 4) returns an integer value in the range [0,4] for the predicted class (see Section 4.1). This integer value c_i is then used as the new input to the following Poisson-style exponential smoothing equation:

$$e_0 = c_0, t = 0$$
(2)
$$e_t = \alpha c_{t-1} + (1 - \alpha) e_{t-1}, t > 0$$

Here, e_t represents the new emotional force. After e_t has been derived according to this formula, it is scaled from range [0,4] to range [-100,100] when applied to the mass in the physical model. The term α represents the smoothing coefficient, which defines the relative weights of the raw classification input c_{t-1} and the most recently calculated emotional force e_{t-1} . For the current research α

was set to 0.6, but future research in this direction might consider experimentally determining the optimal value for this constant.



Figure 4: Diagram displaying the overall process at work, beginning with the speech act from a human agent and resulting in affective movement from the robot.



Figure 5: Emotion detector. (1) Document topic structured built (2) For each sentence, topic probabilities are extracted & used to train the classifier (3) Features are extracted from the utterance by the trained LDA model and (4) presented to the trained classifier to predict emotional state [40]

3.4 Implementation

3.4.1 Emotion Detector Component. We enhanced the three-state sentiment detector described in [40] to infer five states of emotional valence {strong negative, medium negative, neutral, medium positive, strong positive}. Socher et al. [33] reported that five levels of positivity was sufficiently fine-grained to capture the continuous values human evaluators reported and this informed the number of states we chose to detect. Our validation suggests that humans do indeed distinguish the different categorical levels of positivity expressed through the mass-spring component which gives the appearance of smooth transition from one emotion to another.

The mode consists of two processing steps: (i) extract the topic probabilities from each document in the set (items 1 and 3 in Figure 5), and (ii) use these features to predict the emotion valence of individual sentences as yet unseen by the model (items 2 and 4 in Figure 5). Training of the LDA model and the classifier (items 1 and

2 in Figure 5) was done outside of the robotic architecture and the results saved to files. These were subsequently used to initialize the predictor component in the robotic cognitive architecture; in principle, training could also take place in the architecture.

Latent Dirichlet Allocation (LDA) has been shown to be effective for inferring affect in conversational speech [29, 30, 39, 40]. We used the Gensim [26] implementation of LDA as the feature extractor of our model; we used the default hyper-parameter values. The generative model assumes a number of topics over which an initial distribution of documents is estimated. For this implementation, we set the number of topics to be 100, as described in [38] since we used their dataset.

We used a multi-layer perceptron (MLP) with two layers of 50 artificial neurons each; these values were selected based on a parameter sweep using Scikit-learn's GridSearchCV method. The parameter sweep also compared classifiers, i.e., MLP, Support Vector Machine, and Linear Regression; we found this particular MLP configuration to provide the best performance. We trained the model using the *tanh* activation function with a constant learning rate with initial value 0.001 and adaptive moment estimation (i.e., "Adam") as a fast optimizer. We used a stable, widely-used implementation of the MLP classifier from Scikit-learn [22]. We trained the LDA model and the classifier outside of the robotic architecture and the results were saved to files. The Detector component used these files to instantiate the trained model in the robotic cognitive architecture.

3.4.2 Model Training. Training input to the model used the dataset described in [38]. This consisted of individual sentences drawn from 448 documents with an average word count of 258 words, the largest containing 1,732 and the smallest, 2. The documents were constructed from selected interview transcripts from 106 male and female participants with PD, living in the community, who participated in a study [36] which asked them to recall two types of experiences they had during the past week: a frustrating one and an enjoyable one. The robot running the emotion detection model could then be expected to accurately predict emotion from utterances spoken in a similar contextual domain.

Ground truth labels for our model were obtained as follows. Twodimensional (valence, arousal) emotion values for each sentence of the dataset were generated by human evaluators who used a Web-based implementation of the Circumplex model of emotion as described in [38]. In the Circumplex, valence and arousal can range from -100 (most negative/calm) to +100 (most positive/aroused) with 0 considered to be neutral. In that study, the researchers found inter-rater reliability for arousal to be low and therefore used only valence for model training and prediction. The human evaluators (N = 1,058) rated 439 documents of various lengths (269 describing a frustrating experience and 170 describing an enjoyable experience) for a total of 7,713 sentences. Each document was rated by at least four evaluators who rated between two and four documents each depending on the length. We used human evaluators drawn from the general population rather than, for example, asking PD persons themselves to label the data in some assisted way or ask some PD experts to conduct the task. The reason for this is as follows.

Correctly detecting the emotion of PD persons is a challenge because their facial expression do not match what they convey through words. Often, a person with PD can have an angry or apathetic-looking expression even when they are talking about joyful experiences. Research has shown that even specialists have a very hard time inhibiting their incorrect impressions of the person with PD when faced with dissonant emotion expressions across channels [37]. When evaluating the emotion expression in the unaffected channel alone (content of speech), people should have no issues with detecting the correct emotion. This is exactly what our raters did: they read text transcriptions of interviews conducted with people with PD - they never saw the facial masking in conjunction with what was being said.

Using k-means clustering of the (x,y) data we collected from the AMT study, we found five classification center points. This classification gives valence scores of -100 to -10 as negative, -10 to 25 as neutral, and 25 to 100 as positive. Upon model evaluation, we found that classification by constant valence scores of -100 to -24 as negative, -24 to 24 as neutral, and 24 to 100 as positive give the most predictive ranges to use for classification as evaluated by its F_1 score, a common measure of classifier performance. We suggest that users view the Circumplex as a square graph, with the center as true neutral. We therefore use constant valence values away from the center to delineate classification boundaries.

3.4.3 Expressor Component. In order for the mass-spring to operate in real-time as the robot engages in a human interaction, we built a Java component which uses a second order integrator implemented by the Apache Commons Ordinary Differential Equations (ODE) package [11]. Figure 4 displays the control and data flow for this entire process. The component is constantly running in parallel with the other processes in the pipeline. A method within the mass-spring, named Ready_update is called with rapid frequency by the scheduler of the cognitive architecture. Each time it integrates Equation 1 over a given time step, it first updates F to reflect the current force derived from the most recent emotional valence prediction received from the Detect component. Each time a prediction is generated, this force attribute is modified accordingly for six seconds before reverting to 0. It then saves the state of the particle at the end of this time-step, and begins the next time-step from this saved state. This state data includes *x*, the position of the particle, which is used to determine which gesture to send along to the Nao for execution.

As is described in the previous section and displayed in Figure 2, the Nao robot is able to produce seven different gestures, intended to be dispersed evenly along the valence axis from extreme negativity to extreme positivity. In order to translate from the x produced by the integrator into an embodied behavior, we defined numerical thresholds between each gestural space along this axis, derived through observation of the distance from neutral that the particle reached during the application of various positive and negative forces. At each time step, after integration, this component assesses whether the particle has crossed a threshold from one gestural range into another, at which point the robot is instructed to switch from one behavior to another. Each behavior is designed so that it may be repeated continuously for the entire time that the particle occupies the corresponding region. Whenever this position crosses one of the predefined thresholds into a region associated with a different body language behavior, the new desired movement is reported back to the Goal and Action manager component of the

cognitive architecture, which relays it to the Robot Controller Nao component for execution.



Figure 6: Still photo of Nao robot and PD person (i.e., a confederate) used in online evaluation of model.

4 METHODS

4.1 Design and Procedures

To validate our system we conducted three on-line experiments. The procedures were the same for each of the experiments, but the robot emoting was varied as explained below. In each experiment we asked participants to watch a video of a person being accompanied by an emoting robot (see still shot in Figure 6). We conducted the experiments on-line and used videos because we wanted, for a fair comparison, to keep constant what the person was saying across conditions. This would have been problematic in a natural interaction scenario between the participant and the assisted person. Moreover, an in-person interaction presented ethical concerns: using a confederate that actually suffered from PD for testing the robot at this stage of prototyping would have put unnecessary burden on someone vulnerable from a health-perspective, and would have created potential for stigma, while using an actor to impersonate someone with PD mimicking all motor aspects of the disorder would have constituted deception of the participant much beyond what was needed for the purpose of the validation of the system. We thus opted for a video in which a male actor speaks directly to the camera reproducing the facial masking and the affectless tone of voice that is typical of PD. The script was extracted verbatim from an actual interview with a person with PD who was talking about one enjoyable and one frustrating experience they had had the previous week. This interview was set aside from the dataset that was used to train the prediction model.

The emoting robot's behavior was varied across experiments and conditions in the following way: in Experiment 1 we compared the robot emoting based on the model described above (including the Detector and Expressor components) with a video in which the robot was emoting randomly. The purpose of this experiment was to obtain a baseline of how much people associate the robot's gestures with the person's speech. We hypothesized that a higher association to the content of speech would be perceived for the model-based emoting an the random emoting. In Experiment 2 we compared the robot emoting based on the mass-spring model described in this paper, to the robot emoting using the model but without the mass-spring element. The robot produced gestures corresponding to the emotional content of every utterance made by the person. We hypothesized that there would be no differences between the mass-spring and no mass-spring models when the emoting frequency was high (for every utterance). In Experiment 3 we used the same comparison as in Experiment 2, but this time the robot was emoting at a low frequency, only expressing the emotional content for every third utterance made by the person. This experiment was meant to showcase the robustness of the massspring model. We hypothesized that when the frequency of emoting drops, which might happen due to failures of the LVASR or of the Detector component, the mass-spring model would compensate for these errors. By returning to a neutral state instead of perseverating on one particular gesture, it would improve the perceived emoting accuracy. In all experiments participants were given the following instructions: "In this video you will watch a person being interviewed about some enjoyable and some frustrating experiences he's had in the past week. Accompanying him is his assistive robot. Please watch the video carefully. You will be asked questions about the person and his assistive robot."

The study used a between-group design, each participant being randomly assigned to one of the conditions mentioned above. After watching the video, participants answered the following questions: "Is the robot's behavior connected to what is being said?", which participants answered with either "yes" or "no"; and "How well do you feel the robot's gestures matched what was being said?", which participants answered on a five-point Likert scale from "not at all" to "very much so". Additionally, they answered further open-ended and multiple-choice questions about the robot's behavior and the person in the video, but analyses of those answers are beyond the scope of this paper. A total of 161 participants completed the study on Amazon Mechanical Turk and also passed our attention checks (42.6% Female, *Mean* age = 37.4 years, SD = 11.85). The research was approved by the university's Institutional Research Board (IRB), and participants were compensated with \$1.00 USD for their time.

5 RESULTS

5.1 Experiment 1

When the robot was emoting based on our model, 80.7% of the participants indicated that the robot's behavior was connected to what was being said, significantly more than when the robot was emoting randomly, 41%, $\chi^2(2) = 8.86$, p = 0.003 (see Figure 7). Also, when the robot was emoting based on the model, the mean participant rating of how well the robot gestures matched what was being said (*Mean* = 1.96, *SD* = 0.87) was significantly higher than the mean participant ratings when the robot was emoting randomly (*Mean* = 1.14, *SD* = 0.91), t(53) = 3.41, p = 0.001 (see Figure 7). This suggests that people understand the robot's emotive gestures as related to the content of the person's speech.

5.2 Experiment 2

In this high-frequency emoting comparison between our massspring model and the same model without the mass-spring element, we found no differences between the robot emoting based on the mass-spring model and the robot emoting without the massspring element. There was no significant difference between the two conditions in terms of perceived connection between the robot's gestures and what was being said (mass-spring:76%, no mass-spring: 74%, $\chi^2(2) = 0.03$, p = 0.868), or between the mean participant ratings of how well the robot's gestures matched what was being said (mass spring: *Mean* = 1.92, *SD* = 1.04, no mass-spring: *Mean* = 1.74, *SD* = 1.10), t(46) = 0.58, p = 0.559.

5.3 Experiment 3

When the frequency of emoting was low however, the robot emoting using the mass-spring model outperformed the robot emoting without the mass-spring element. The use of the mass-spring led to a higher perceived connection between the robot's gestures and what was being said (mass-spring: 85% no mass-spring: 48%, $\chi^2(2) = 8.65$, p = 0.003) and participants rated the mass-spring emoting as better matching what was being said (mass spring: Mean = 1.96, SD = 1.09, no mass-spring: Mean = 1.16, SD = 0.900) t(56) = 3.07, p = 0.003. This suggests that the mass-spring emoting model is robust to potential LVASR or recognition errors and is perceived by observers to be more accurate at expressing emotions from speech content.



Figure 7: Ratings of perceived robot behavior in (A) highfrequency and (B) low-frequency emoting conditions. Top and bottom rows indicate participant responses to questions (1) and (2).

6 **DISCUSSION**

We hypothesized that gestures could be an effective mode for conveying emotion in a robot, and that the mass-spring would be a robust design that ensures high perceived accuracy of emoting even when emotions are detected at low frequency. This might occur, for example, when utterances are spatially separated because of pauses in speech or when detection frequency is reduced to conserve power in an embedded system. We further recognized that emotion detection models are fallible and that some means would be necessary to smooth the impact of erroneous predictions on the robot's expression or when critical, emotion-bearing utterances are omitted or not recognized by the LVASR component. Therefore, we designed the mass-spring component to mediate between the prediction component and the gesture generation to serve this purpose.

Our findings support our hypotheses that humans do indeed distinguish the different categorical levels of positivity expressed through the mass-spring component which gives the appearance of smooth transition from one emotion to another. Compared to random emoting, when evaluating the overall perception of the robot's behavior and gesticulations produced by the detector and massspring components in concert, participants perceived a significantly higher association between the robot's emoting and the person's content of speech. The findings suggest that this system provides a suitable basis for a emoting robot companion for persons living with Parkinson's disease. Additionally, the mass-spring dynamical system led to greater perceived association between the robot's gestures and the emotional content of the speaker's utterances than a model without a mass-spring when the emoting was done at a low-frequency (for every third utterance). This suggests that the mass-spring provides necessary robustness for the emoting system, which is critical given the targeted population - people with PD often have difficulties with speech production, which might make the error rate of the LVASR component higher. Our results suggest that the mass-spring could help compensate for this.

During in-lab system testing, we noticed some delay, not exceeding 1.5 s on average between what the person uttered and the generated gesture depending on how the speech end-point was detected by the ASR component. This also impacted the accuracy of the predicted emotion. If there were disfluencies in the speech, or the person paused, the utterances would be broken up into smaller segments. While this reduced the delay, this increased the likelihood of there being insufficient context to accurately detect the intended emotion. As noted in Section 22, the LDA method used in the detector has been shown to be more accurate than, for example, a sentiment lexicon approach such as LIWC in detecting emotion in short bursts of text. When evaluating the gesture set, we found that while overall the participants were able to recognize the relative positivity of the set, within the three less positive and three more positive sets surrounding the neutral gesture, there was some lack of distinction. We attribute this to having insufficient context to situate the gestures. Further research is needed to design a more definitive sequence.

At this point we integrated a pre-trained speech recognizer [41]. Since speech recognition is not within the scope of this particular paper, one way in which we circumvent noise and multiple people is by using a microphone. However, we recognize that for this to be completely functional, speech recognition in noisy environments is essential. We are however showing that the spring-model actually aids with noisy or low-frequency data and performs better than a model without a spring component.

6.1 Future Work

We based our model on Russell's Circumplex bi-axial model of emotion. It cannot be easily flattened onto a single dimension while retaining specific, discrete emotional categories because multiple recognizable emotions may exist at the same general valence while deviating in arousal. However, in future work it would be interesting to compare performance of the system on a continuous vs. a discrete model of emotion.

A future study would do well to investigate the contribution of each component of the full system in an empirical experiment under varying conditions of emotion detection frequency and error rates. Utterances containing high and low emotional variance and sequences of abrupt transitions between positive and negative emotions would further help in the analysis of mass-spring's contribution. Furthermore, human emotional state can change in far more complex ways and in more subtle gradations than the five emotional categories detected in this system. Refining the emotion detector to generate not only valence but arousal measures would could reproduce more accurately the complexity of human emotion. The challenge, then, will be to design suitable emotion expressions in the robot that reflect this complexity.

7 CONCLUSION

We developed and evaluated a model which detected five degrees of emotional valence in the continuous speech of a person with PD and used a spatial-temporal dynamical system to compensate for emotion detection errors and frequency of emission. We embedded the model in the DIARC robotic cognitive architecture running in a Nao robot which expressed emotion using its body movement. The system equips the robot with the ability to provide immediate feedback on the emotional state of the person with PD during conversations with their care-givers or in social-situations. Prior research has shown providing feedback on the emotion content of a conversation is not only beneficial for improving the social interaction with the PD patient, it can improve the quality of life in the home. Since human emotion is communicated via multiple modalities, and through different channels, (e.g., voice, facial expressions, gestures) situating such a tool in a robot that appropriately controls its expressive motors could compensate for the compromised vocal and facial modalities when communicating emotion.

We found encouraging results that showed participants in our study connected the robot's gestures to what was being said. Once enhanced with finer prediction resolution and further tuning of the dynamical system, the robot should be able to express emotion using any bodily movement available in a natural way and under a variety of conditions (e.g., noisy speech, rapid emotional changes). We envision that this system can be generalized to serve as a conversational agent which can monitor the emotional content between any two individuals and provide immediate feedback on the emotion content during the course of the conversation.

ACKNOWLEDGMENTS

The work is supported by the National Science Foundation under Grant No.: IIS-1316809.

REFERENCES

 2015. Four Teams Win IARPA's ASpIRE Challenge. https://www.afcea.org/ content/?q=Blog-four-teams-win-iarpas-aspire-challenge. Accessed: 2019-11-12.

- [2] Kenji Amaya, Armin Bruderlin, and Tom Calvert. 2000. Emotion from Motion. Proc. Graphics Interface Conf (12 2000).
- [3] E. Cambria. 2016. Affective Computing and Sentiment Analysis. *IEEE Intelligent Systems* 31, 2 (Mar 2016), 102–107. https://doi.org/10.1109/MIS.2016.31
- [4] Pauline Chevalier, Brice Isableu, Jean-Claude Martin, and Adriana Tapus. 2016. Individuals with autism: Analysis of the first interaction with nao robot based on their proprioceptive and kinematic profiles. In Advances in robot design and intelligent control. Springer, 225–233.
- [5] Ze-Jing Chuang and Chung-Hsien Wu. 2004. Multi-modal emotion recognition from speech and text. International Journal of Computational Linguistics & Chinese Language Processing, Volume 9, Number 2, August 2004: Special Issue on New Trends of Speech and Language Processing 9, 2 (2004), 45–62.
- [6] B. de Gelder, A.W. de Borst, and R. Watson. [n.d.]. The Perception of Emotion in Body Expressions. Wiley Interdisciplinary Reviews: Cognitive Science 6, 2 ([n.d.]), 149–158. https://doi.org/10.1002/wcs.1335
- [7] P. Ravindra De Silva and Nadia Bianchi-Berthouze. 2004. Modeling human affective postures: an information theoretic characterization of posture features. *Computer Animation and Virtual Worlds* 15, 3-4 (2004), 269–276. https://doi.org/ 10.1002/cav.29
- [8] Paul Ekman. 1992. An argument for basic emotions. Cognition and Emotion 6, 3-4 (1992), 169–200. https://doi.org/10.1080/02699939208411068
- [9] Sarah Fdili Alaoui, Cyrille Henry, and Christian Jacquemin. 2014. Physical modelling for interactive installations and the performing arts. *International Journal of Performance Arts and Digital Media* 10, 2 (2014), 159–178.
- [10] Huanghao Feng, Anibal Gutierrez, Jun Zhang, and Mohammad H Mahoor. 2013. Can NAO robot improve eye-gaze attention of children with high functioning autism?. In 2013 IEEE International Conference on Healthcare Informatics. IEEE, 484–484.
- The Apache Software Foundation. 2019. Ordinary Differential Equations Integration. https://commons.apache.org/proper/commons-math/userguide/ode.html. Accessed: 2019-08-15.
- [12] Taabish Gulzar, Anand Singh, Dinesh Kumar Rajoriya, and Najma Farooq. 2014. A systematic analysis of automatic speech recognition: an overview. Int. J. Curr. Eng. Technol 4, 3 (2014), 1664–1675.
- [13] M. Han, C. Lin, and K. Song. 2013. Robotic Emotional Expression Generation Based on Mood Transition and Personality Model. *IEEE Transactions on Cybernetics* 43, 4 (Aug 2013), 1290–1303. https://doi.org/10.1109/TSMCB.2012.2228851
- [14] M. Harper. 2015. The Automatic Speech recogition In Reverberant Environments (ASpIRE) challenge. In 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). 547–554. https://doi.org/10.1109/ASRU.2015.7404843
- [15] Mary Harper. 2015. The automatic speech recogition in reverberant environments (ASpIRE) challenge. In 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU). IEEE, 547–554.
- [16] C. Hutto and Eric Gilbert. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. https://www.aaai.org/ocs/index.php/ ICWSM/ICWSM14/paper/view/8109. In International AAAI Conference on Web and Social Media.
- [17] Luthffi Idzhar Ismail, Syamimi Shamsudin, Hanafiah Yussof, Fazah Akhtar Hanapiah, and Nur Ismarrubie Zahari. 2012. Robot-based intervention program for autistic children with humanoid robot NAO: initial response in stereotyped behavior. *Procedia Engineering* 41 (2012), 1441–1447.
- [18] Derek McColl and Goldie Nejat. 2014. Recognizing Emotional Body Language Displayed by a Human-like Social Robot. *International Journal of Social Robotics* 6 (04 2014). https://doi.org/10.1007/s12369-013-0226-7
- [19] E. Mower, M. J. Mataric, and S. Narayanan. 2011. A Framework for Automatic Human Emotion Classification Using Emotion Profiles. *IEEE Transactions on Audio, Speech, and Language Processing* 19, 5 (July 2011), 1057–1070. https: //doi.org/10.1109/TASL.2010.2076804
- [20] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs Up?: Sentiment Classification Using Machine Learning Techniques. In Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10 (EMNLP '02). Association for Computational Linguistics, Stroudsburg, PA, USA, 79-86. https://doi.org/10.3115/1118693.1118704
- [21] S.-T Park, Lilia Moshkina, and Ronald Arkin. 2010. Recognizing Nonverbal Affective Behavior in Humanoid Robots. Intelligent Autonomous Systems 11, IAS 2010 (01 2010). https://doi.org/10.3233/978-1-60750-613-3-12
- [22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [23] J. W. Pennebaker, R. L. Boyd, K. Jordan, and K. Blackburn. 2015. The development and psychometric properties of LIWC2015. University of Texas at Austin, Austin, TX. https://doi.org/10.15781/T29G6Z
- [24] Robert Plutchik. 1980. A general psychoevolutionary theory of emotion. In Theories of emotion. Elsevier, 3-33.
- [25] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz,

Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. The Kaldi Speech Recognition Toolkit. In *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding* (Hilton Waikoloa Village, Big Island, Hawaii, US). IEEE Signal Processing Society. IEEE Catalog No.: CFP11SRW-USB.

- [26] Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. http://is.muni.cz/publication/884893/en. In Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks. ELRA, Valletta, Malta, 45–50.
- [27] James A Russell. 1980. A circumplex model of affect. Journal of personality and social psychology 39, 6 (1980), 1161.
- [28] Matthias Scheutz, Thomas Williams, Evan Krause, Bradley Oosterveld, Vasanth Sarathy, and Tyler Frasca. 2019. An overview of the distributed integrated cognition affect and reflection DIARC architecture. In *Cognitive Architectures*. Springer, 165–193.
- [29] Mohit Shah, Chaitali Chakrabarti, and Andreas Spanias. 2015. Within and crosscorpus speech emotion recognition using latent topic model-based features. EURASIP Journal on Audio, Speech, and Music Processing 2015, 1 (25 Jan 2015), 4. https://doi.org/10.1186/s13636-014-0049-y
- [30] M. Shah, L. Miao, C. Chakrabarti, and A. Spanias. 2013. A speech emotion recognition framework based on latent Dirichlet allocation: Algorithm and FPGA implementation. In 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. 2553–2557. https://doi.org/10.1109/ICASSP.2013.6638116
- [31] Syamimi Shamsuddin, Hanafiah Yussof, Luthffi Ismail, Fazah Akhtar Hanapiah, Salina Mohamed, Hanizah Ali Piah, and Nur Ismarrubie Zahari. 2012. Initial response of autistic children in human-robot interaction therapy with humanoid robot NAO. In 2012 IEEE 8th International Colloquium on Signal Processing and its Applications. IEEE, 188–193.
- [32] C. Shan, S. Gong, and P. W. McOwan. 2007. Beyond Facial Expressions: Learning Human Emotion from Body Gestures. In Proc. BMVC. 43.1–43.10. https://doi. org/10.5244/C.21.43
- [33] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference* on empirical methods in natural language processing. 1631–1642.
- on empirical methods in natural language processing. 1631–1642.
 [34] Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. arXiv preprint arXiv:1903.09588 (2019).
- [35] Yla R. Tausczik and James W. Pennebaker. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. https://doi.org/10.1177/ 0261927X09351676. Journal of Language and Social Psychology 29, 1 (2010), 24–54. https://doi.org/10.1177/0261927X09351676
- [36] L. Tickle-Degnen, T.D. Ellis, M. Saint-Hilaire, C. Thomas, and R. C. Wagenaar. 2010. Self-management rehabilitation and health-related quality of life in Parkinson's disease: A randomized controlled trial. *Movement Disorders* 25 (2010), 194–204.
- [37] Linda Tickle-Degnen and Kathleen Doyle Lyons. 2004. Practitioners' impressions of patients with Parkinson's disease: the social ecology of the expressive mask. *Social Science & Medicine* 58, 3 (2004), 603–614.
- [38] Andrew P. Valenti, Meia Chita-Tegmark, Michael Gold, Theresa Law, and Matthias Scheutz. 2019. In Their Own Words: A Companion Robot for Detecting the Emotional State of Persons with Parkinson's Disease. In 11th International Conference on Social Robotics. Springer, Madrid, Spain, 1–10.
- [39] Andrew P. Valenti, Meia Chita-Tegmark, Theresa Law, Alexander W. Bock, Bradley Oosterveld, and Matthias Scheutz. 2019. Using topic modeling to infer emotional state: When your face and tone of voice don't say it all: Inferring emotional state from word semantics and conversational topics. In Workshop on Cognitive Architectures for HRI: Embodied Models of Situated Natural Language Interactions at AAMAS 2019. Montreal, Canada.
- [40] Andrew P. Valenti, Meia Chita-Tegmark, Linda Tickle-Degnen, Alexander W. Bock, and Matthias J. Scheutz. 2019. Using topic modeling to infer the emotional state of people living with Parkinson's disease. Assistive Technology (2019), 1–10. https://doi.org/10.1080/10400435.2019.1623342
- [41] Krisztián Varga. 2017. Kaldi ASR: Extending the ASpIRE model. https://chrisearch.wordpress.com/2017/03/11/speech-recognition-using-kaldiextending-and-using-the-aspire-model/. [Online; accessed 30-August-2019].
- [42] JD Velsquez. 1997. Modeling emotions and other motivations in synthetic agents. *Aaai/iaai* (1997), 10–15.
- [43] P. Xiaolan, X. Lun, L. Xin, and W. Zhiliang. 2013. Emotional state transition model based on stimulus and personality characteristics. *China Communications* 10, 6 (June 2013), 146–155.
- [44] Guo L. Yang, Jin H. Zhang, and Hui Sun. 2012. Design of Emotional Interaction System Based on Affective Computing Model. Applied Mechanics and Materials 198-199 (09 2012), 367. https://login.ezproxy.library.tufts.edu/login?url=https: //search.proquest.com/docview/1443259693?accountid=14434 Copyright - Copyright Trans Tech Publications Ltd. Sep 2012; Last updated - 2018-10-05.