ECA Control using a Single Affective User Dimension

Fred Charles
Teesside University
School of Computing
Middlesbrough, UK
f.charles@tees.ac.uk

Florian Pecune
CNRS - LTCI
Telecom ParisTech
Paris, France
pecune@enst.fr

Gabor Aranyi
Teesside University
School of Computing
Middlesbrough, UK
g.aranyi@tees.ac.uk

Catherine Pelachaud*
CNRS - LTCI
Telecom ParisTech
Paris, France
pelachaud@enst.fr

Marc Cavazza*
Teesside University
School of Computing
Middlesbrough, UK
m.o.cavazza@tees.ac.uk

ABSTRACT

User interaction with Embodied Conversational Agents (ECA) should involve a significant affective component to achieve realism in communication. This aspect has been studied through different frameworks describing the relationship between user and ECA, for instance alignment, rapport and empathy. We conducted an experiment to explore how an ECA's non-verbal expression can be controlled to respond to a single affective dimension generated by users as input. Our system is based on the mapping of a high-level affective dimension, approach/avoidance, onto a new ECA control mechanism in which Action Units (AU) are activated through a neural network. Since 'approach' has been associated to prefrontal cortex activation, we use a measure of prefrontal cortex left-asymmetry through fNIRS as a single input signal representing the user's attitude towards the ECA. We carried out the experiment with 10 subjects, who have been instructed to express a positive mental attitude towards the ECA. In return, the ECA facial expression would reflect the perceived attitude under a neurofeedback paradigm. Our results suggest that users are able to successfully interact with the ECA and perceive its response as consistent and realistic, both in terms of ECA responsiveness and in terms of relevance of facial expressions. From a system perspective, the empirical calibration of the network supports a progressive recruitment of various AUs, which provides a principled description of the ECA response and its intensity. Our findings suggest that complex ECA facial expressions can be successfully aligned with one high-level affective dimension. Furthermore, this use of a single dimension as input could support experiments in the fine-tuning of AU activation or their personalization to user preferred modalities.

*corresponding authors

ICMI 2015, November 9–13, 2015, Seattle, WA, USA.

DOI: http://dx.doi.org/10.1145/2818346.2820730.

Categories and Subject Descriptors

H.5.1 [**Multimedia Information Systems**]: Artificial, Augmented, and Virtual Realities; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces

Keywords

Embodied Conversational Agent; Brain-Computer Interface; Functional Near-Infrared Spectroscopy (fNIRS)

1. INTRODUCTION AND RATIONALE

One important research direction to improve ECA use in user interfaces is to study the fundamental properties of the ECA-user relationship. These have been described through various concepts such as *empathy* and *alignment*. For instance, ECA are able to induce *empathy* in a human user, which usually translates into exploring appropriate ways for the agent to reveal its internal state. Furthermore, *alignment* aims at mapping between the recognition of facial expressions, in terms of emotion recognition, of the user and the elicitation of the emotional response from the virtual agent through the generation of appropriate facial expressions.

Here, we investigate the potential of using a single high-level emotional dimension as an input to control a virtual agent's response. We aim at using a Brain-Computer Interface (BCI) in a Neurofeedback (NF) setting to provide the user with compelling multimodal visual feedback using complex generation of appropriate facial expressions in response to their (perceived) disposition towards the virtual agent.

In this paper, the ECA is designed to provide alignment between the users' positive thoughts, collected via fNIRS-based BCI input signals, and the generation of perceived internal emotions of a virtual agent through the display of appropriate facial expressions. In the next sections, we provide an overview of the related works in virtual agents, and their ability to generate emotions, as well as a detailed description of the virtual agent configuration and the experimental setup we have designed. Finally, we present the results of our experiments.

2. PREVIOUS AND RELATED WORK

To enhance human-agent interaction, several studies have reported the importance of an agent's ability to display emo-

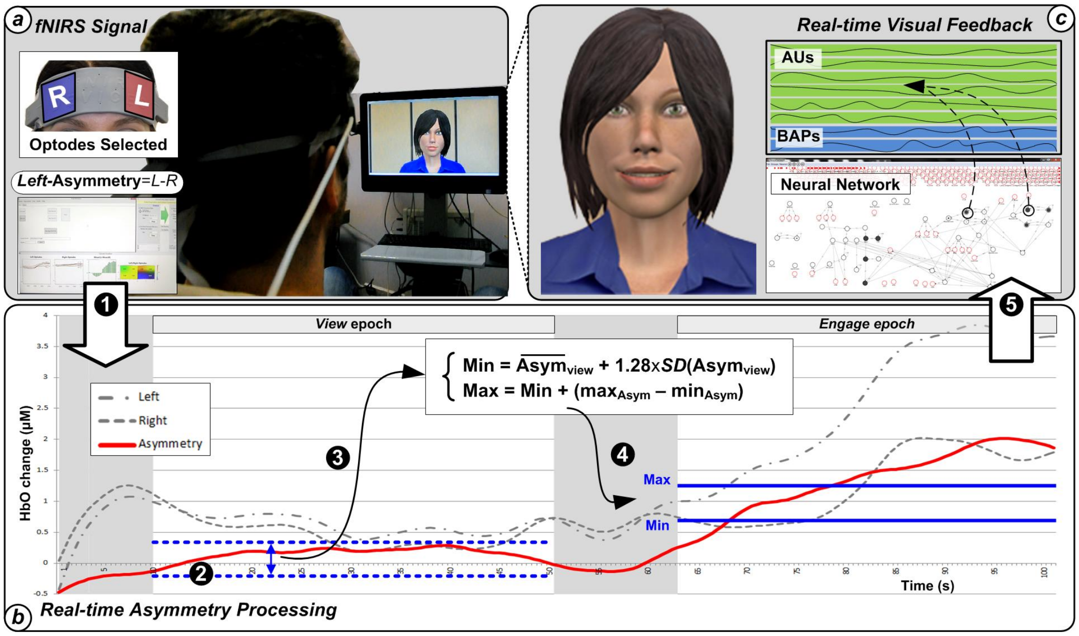

Figure 1: System Overview. Brain signals are collected through fNIRS system ⓐ, where left-most and right-most channels are processed to generate a left-asymmetry score ①. During the *View* epoch ②, the left-asymmetry values are used to define the *Min* and *Max* bounds ③ to be used during the *Engage* epoch where the real-time left-asymmetry scores ④ are normalized before being used as single input ⑤ to the virtual agent's facial expressions (AUs) and body parameters (BAPs) under the neural network's control ⓒ.

tions [5, 25, 26]. However, the appropriateness of the displayed emotions is paramount. Expressing an emotion that is inappropriate in the context of the interaction will damage the interaction more than not showing any emotion at all [25]. Moreover, *empathic* virtual agents have a greater impact on users' perception than agents that are fully impulsive and express their felt emotions without consideration of their interlocutors' emotions nor the social context of the interaction [5, 25]. Previous research has investigated both empathic agents supporting the user in learning situations [26] and agents that elicit empathy from users. The use of physiological signals has been experimented in both contexts [16, 26]. Gilroy et al. [16] have described the use of EEG prefrontal asymmetry to capture empathy towards a virtual agent during an interactive narrative, but have not explored the agent's expressions, instead using a simplified feedback mechanism based only on color saturation.

Agents display their emotional states through their verbal and non-verbal behaviors. While earlier models focused on the six prototypical expressions of emotion, latest models allow modeling a large variety of facial expressions as a combination of the expressions of emotions [1] as sequences of multimodal signals [22], or even as a blend of expressions [23]. Models may rely on pure combinatorial approaches where the expression of an emotion is the result of an algebraic combination of other expressions [1] of regions on the face [7, 23]; other models have relied on corpora annotation [22] where videos are carefully annotated to extract multimodal signals expressing emotions. Others have applied a perceptual approach [15, 17] where human users are asked to create expressions of the virtual agents for given emotions. Yet other approaches have implemented the dynamic evolution of facial expression of emotion following Scherer's appraisal model [2, 11].

Additionally, an important research topic in relation to the idea of 'alignment' between human users and virtual agents is 'rapport' defined in [18] as the feeling of being 'in-sync' with a conversational partner. They also report that virtual agents can create rapport during interactions with human users by generating proper verbal, and non-verbal behaviors. This provides a setting whereby we assume that generated non-verbal behaviors (i.e., facial expressions) are the reflection of the human user which thus creates a relationship between the user's input signals and the resulting visual feedback generated by the virtual agent. Lately, Cassell and colleagues [32] have taken into account the expressed and perceived mental state of each participant of a dyad to model 'rapport'. Their model reasons about the beliefs and intentions of one participant and how it perceives its interlocutor's state of mind.

In the next sections, we present the overall system design, introducing how user disposition towards the agent can be captured by a fNIRS-based BCI and mapped onto control parameters for the ECA.

3. SYSTEM OVERVIEW

We designed a software platform to experiment with ECA control from a single affective dimension, which is captured through an fNIRS BCI (see Figure 1). From the user's perspective, the ECA behaves as an autonomous agent, responding to what it perceives as the user's mental disposition towards itself. The ECA possesses a range of spontaneous and expressive behaviors, ranging from idle animations, with (neutral) spontaneous motion and postural changes, to affective expression of graded intensity. Users are instructed to express positive feelings towards the agent in order to capture its interest. This should in turn result in the ECA responding to the user with an expression matching the perceived interest in both valence and intensity. Users' attitudes are captured through their levels of Pre-Frontal Cortex (PFC) asymmetry, using fNIRS to measure differences in activity in *left* versus *right* PFC (Figure 1(a)). This is based on extensive work relating the approach/withdrawal affective dimension to PFC asymmetry, measured through EEG [12] as well as MRI [33]. We have adopted fNIRS because of its lesser sensitivity to artefacts compared to EEG and the fact that it is well-suited to capture activity in the PFC, in particular its dorsolateral region [13]. Furthermore, experience gained when studying PFC asymmetry through fMRI [9] could provide guidance to design fNIRS experiments. From the ECA perspective, the main challenge is to design an appropriate control process that would interpret and respond appropriately to the level of approach expressed mentally by the user. The system uses a network-like control representation to relate a single input to an array of Action Units (AU) and Body Action Parameters (BAP) that provide low-level control for the ECA's animation (Figure 1(c)).

An overview of the operation of the system is as follows. The system captures PFC asymmetry in real-time and produces a matching ECA response. The PFC asymmetry baseline, which is subject-dependent, is acquired through an epoch during which subjects watch the ECA (displaying a neutral attitude) while carrying out a simple mental counting task (Figure 1(b)). During active engagement by subjects, PFC asymmetry is measured and its increase over the baseline is interpreted as the intensity of the approach towards the ECA. Finally, the value is passed to the control network, which will generate matching ECA behavior, interrupting idle behavior and producing appropriate facial expressions.

In the next sections, we describe in more detail the various components of the system, as well as the techniques for data acquisition and calibration, which have been used to map BCI data onto the ECA controller.

3.1 Affective BCI Input

Previous uses of fNIRS in affective research [21, 31] have investigated emotional responses (to music or video clips) rather than the active expression of emotions as an input modality. Here, we used fNIRS as an active, affective BCI based on fNIRS (see [8] for a general description of fNIRS).

For the purpose of providing real-time NF, we used measurements of changes in oxygenated hemoglobin (HbO) concentration[1], because it is commonly associated with neural activity [28]. We used an fNIR400 Optical Brain Imaging Station by Biopac Systems, with a 16-channel sensor with fixed 2.5cm source-detector separation (Figure 1(a). Raw data and oxygenation values were collected with a 2Hz sampling rate using software provided by the device manufacturer (COBI Studio and FnirSoft v3.5), and was sent to a bespoke remote client experimental software over TCP/IP (using FnirSoft DAQ Tools). To characterize left PFC asymmetry for real-time feedback, we derived a metric of inter-hemispheric difference in HbO change by averaging HbO values over the four leftmost and four rightmost channels (located over the left and right dorsolateral prefrontal cortex, respectively), then subtracting average *Right* from average *Left* [13].

The resulting asymmetry metric (henceforth called *left-asymmetry*) reflects the inter-hemispheric difference in HbO change in micromolar units (μM/L) which defines the single affective dimension as continuous input signal to control the ECA.

[1]As opposed to deoxygenated or total hemoglobin (HbR and HbT, respectively).

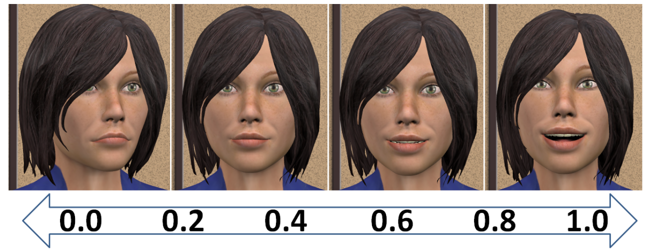

Figure 2: Evolution of the agent's behaviors according to the user's level of engagement.

3.2 Virtual Agent Control Mechanism

Our virtual agent [6] is a synthetic computer animated character designed to autonomously interact with humans in various activities. The agent relies on procedural models involving upper body modalities: facial expression (eyebrow, eyelid, cheek, lip and jaw), gaze, head and upper torso motions. The facial expressions of the virtual agent are defined using FACS, Facial Action Coding System [14], where face movements are decomposed into several Action Units (AUs). For instance, a raise of inner eyebrows corresponds to AU1, while a jaw dropping corresponds to AU26. Each AU is represented by a value varying from 0 (no activation) to 1 (maximum activation). For the body animation, our agent follows the MPEG-4 standard which defines several Body Animation Parameters (BAPs). Each BAP represents the rotation angle of one body joint on one axis.

The ECA responds by aligning its nonverbal behavior to the user's detected empathy level. The detected empathy level is computed over one single affective dimension, ranging from zero to one. As just explained, a very low value of the affective dimension corresponds to a user with very low level of empathy; while a very high value corresponds to a user with very high level of empathy. To achieve agent's behavior alignment with user's detected empathy level, we mapped user's low level of empathy with agent's disengagement and user's high level of empathy with agent's full engagement. Thus, as user's positive thoughts increases, the agent gives more positive feedbacks as sign of its engagement level. For intermediate levels of engagement, the agent displays different variation of positive expressions. The design of the agent's nonverbal behavior, mapped to AUs/BAPs, is drawn from literature [19] and is validated from consensus of a small number of participants.

We consider four different levels of empathy (see Figure 2). If the user's detected level of empathy is minimal (i.e. $\leq$0.2), the agent looks disengaged mirroring to the user his detected mental state. In this case, the agent does not direct its gaze towards the user; it tilts its head and its body, and it displays boredom through a lip pout. Then, whilst the user's detected level of empathy remains below average (i.e.

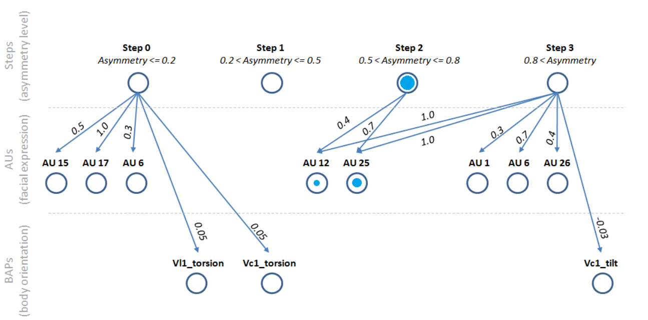

Figure 3: Example of network mapping asymmetry levels to AUs and BAPs. Here the transition step based on the user's perceived level of empathy leads to the activation of AU12 (Lip Corner Puller) and AU25 (Lips Part).

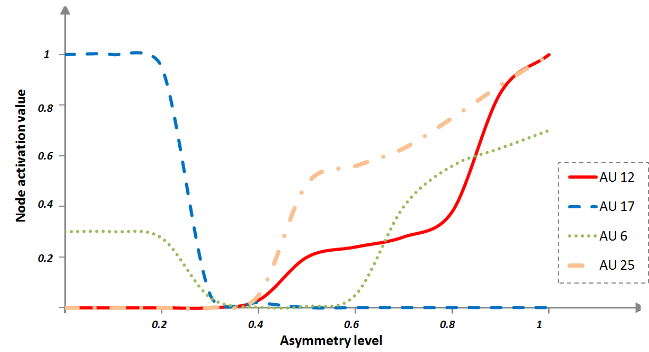

Figure 4: Activation patterns of AUs defined in the network according to the changes in the user's asymmetry levels.

$\leq$0.5), the agent shows a neutral facial expression with its gaze and body oriented towards the user. When the user's detected level of empathy reaches above average ($>$0.5 but $\leq$0.8) the agent shows a more positive attitude, expressing happiness. It keeps its head and its gaze directed at the user, and it displays a smile. Finally, once the user's detected level of empathy reaches near maximum (i.e., $>$0.8), the agent displays a real delight to interact with the user. Its smile widens, crow's feet appear around the eyes as the cheeks are raised and its head slightly tilts backwards.

To provide continuous feedbacks to the user, the agent varies its behavior in real-time. Thus, in order to modify the generated animation in real-time, we developed a graphical tool allowing us to map any kind of input (here, the user's perceived level of empathy mapped from the left-asymmetry value) into animation parameters for our agent (AUs and BAPs). We created a graph where the nodes are the inputs and outputs and the arcs represent the links between the input nodes and the output ones. Like in neural networks, nodes have their own activation values, while arcs have their own weights.

In the graph shown in Figure 3, each of the transition step is represented by an input node whose value ranges from 0 (no activation) to 1 (maximum activation). Depending on the user's perceived level of empathy, a particular node is activated. On the output side, every AU and BAP is also represented by its own node. The values of a node for each AU vary from 0 (no activation) to 1 (maximum activation). So for example, if the node AU12 (Lip Corner Puller) is fully activated (its value is 1), the agent will smile with the greatest intensity. For the BAPs, nodes values range from -1 to 1. These values are then transformed into radian angles to be interpreted by our animation system, as specified by the MPEG-4 norm. For instance, if the value of the node vc1 tilt (cerebral vertebra along the neck) is set to -1, the agent's head will be tilted backwards with an angle of $-\pi$.

Finally, arcs are drawn to connect inputs (left-asymmetry level nodes) to the outputs (AUs and BAPs nodes). For each level of empathy, we define a combination of AUs and BAPs conveying the associated behavior. It corresponds to defining 4 different nonverbal behaviors to be displayed by the virtual agent in response to user's level of left-asymmetry. For instance, to express disengagement (reflecting a low level of left-asymmetry), AU6 (Cheek Raiser), AU15 (Lip Corner Depressor) and AU17 (Chin Raiser) are activated, as well as BAPs vc1 tilt (cerebral vertebra along the neck) and vl1 tilt (lumber vertebra of the trunk). This expression corresponds to a pout with the agent not gazing at the participant. Thus, five arcs are connecting the input node Step0 to the five corresponding output nodes. The weight of each arc connecting level p with p+1 is set by calibrating AUs and BAPs with their extreme values of step p, in order to obtain a smooth animation during the transition phase between step p and step p+1, allowing the agent to display believable behavior.

Each arc has its own weight, representing how much the input node will influence the output one. Weights vary from -1 (negative influence) to 1 (positive influence). The output nodes activation value is computed based on McCulloch and Pitts activation function [20]. The final output activation value is a weighted sum of the inputs where w_{ij} represents the weights of the links and a_i represents the activation values of the input nodes. θ_j represents the initial activation value of the output node.

$$a_j = \sum_{i}^{n} w_{ij} * a_i + \theta_j$$

For instance, if the input node Step0 is fully activated and connected to the output node AU62 (Eyes turn right) by an arc whose weight is 0.3, the agent will slightly look to its right. Therefore, whenever one input node is activated, the corresponding AUs and BAPs are activated. To ensure smooth transitions between the different behaviors associated to each step, we interpolate the animations. When the user's perceived level of empathy generates a change from

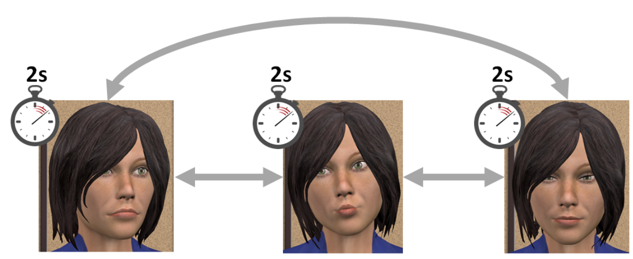

Figure 5: Example of variations of 'idle' behaviors (randomly generated every 2s).

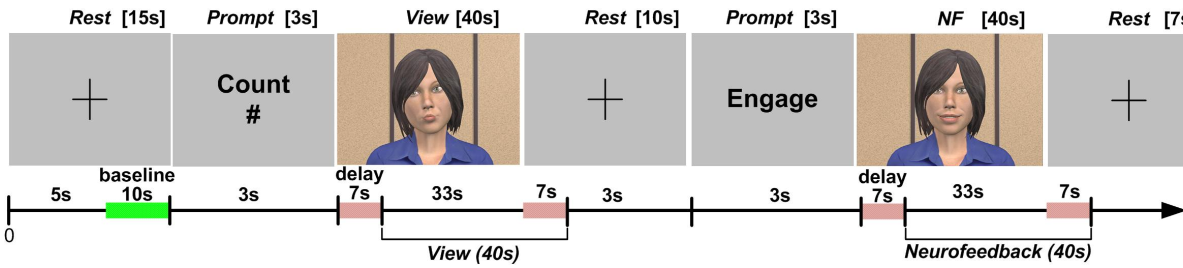

Figure 6: Protocol design, including setting the baseline of the fNIRS system, as well as windowing the data collection during the *View* and *Engage* epochs to account for the delay in hemodynamic response.

one step to another, the corresponding input nodes are not directly activated or inhibited, but instead, the activation and inhibition of these nodes follow an exponential function illustrated in Figure 4.

Furthermore, to avoid the agent displaying repetitively the same facial expression, we introduce alternative combinations of AUs and BAPs for each step. When communicating, different behaviors can be used while conveying similar meanings [24]. The new expressions defined for each step are built as variations of behaviors. Variations can be a change of AU or BAP intensity, a change of AUs (e.g. a change between AU25 (Lips Part) and AU26 (Jaw Drop)), etc. The variation of expressions ensures the agent not to display repetitive behaviors but also not to freeze if remaining in a same step for a long duration. Without such a variation, the agent would display a unique expression that would remain static and thus would appear uncanny. Moreover, it also allows giving continuous feedback to the user about his level of engagement, thus whilst the user is perceived to maintain a level of empathy, a particular, though constantly relevant, behavior should be generated. Hence, we introduce a timer that triggers the display of one of the variants of the behavior of the corresponding level. Empirically, we defined the timer to be 2s for which the user maintains a particular level of left-asymmetry (see Figure 5). This mechanism preserves the naturalness of facial expressions as the agent displays continuously expressive and dynamic behavior.

3.3 Experimental Task

For the design of our BCI experiment to investigate ECA control from a single affective dimension, we followed the recommendation of Solovey et al. [30] for the use of fNIRS in HCI settings. The experimental task consisted in completing 8 identical blocks (preceded by a practice block which was not analyzed). The structure of the blocks is presented in Figure 6. Each block included three epochs: *Rest*, *View*, and *Engage*.

During *Rest* epochs, subjects were instructed to look at a crosshair located in the center of a gray screen to try to clear their head of thoughts and relax. During *View* epochs, subjects were instructed to keep looking at the agent while carrying out a simple mental counting task (counting backwards from 500 by increments of a given integer). This task was included to control for unwanted mental processes (see [29]). During *Engage* epochs, subjects were instructed to engage with the ECA through positive thinking, and to 'cheer her up' with their thoughts. We were deliberately vague with support instructions in order to allow subjects to develop their own cognitive strategies. After completing each block, subjects were asked to describe their strategies in general terms.

During *Engage* epochs, subjects received real-time feedback of their left-asymmetry. To ensure consistent mapping of individual variations in left-asymmetry onto the feedback signal, we used the range of variation of HbO asymmetry during the *View* epoch in each block to determine the mapping of the level of engagement from the user to the visual feedback signal (i.e., ECA's variations in facial expressions - see Figure 1(3)). This was calculated by the experimental software during the last 3s of the *Rest* epoch between the *View* and *Engage* epochs.

We defined the minimum point for mapping *Min* as the mean of left-asymmetry values during the *View* epoch plus 1.28 times their standard deviation. In normally distributed asymmetry scores, this threshold would result in no feedback for 90% of the spontaneous asymmetry variations during the reference (*View*) epoch[2]. To determine the maximum *Max* point for mapping, we increased the threshold asymmetry value for feedback *Min* by the variation range of asymmetry values during the *View* epoch. Asymmetry values within the range [*Min*;*Max*] during the *Engage* epoch were mapped linearly onto the ECA's facial expression, with the same 2Hz frequency as the acquisition of asymmetry values.

Ten subjects participated in the experiment (5 female, Mean age=43.89 years, SD=8.15, Range=[31; 53]). Subjects were right-handed and had no history of psychiatric conditions. They provided written consent prior to participation and received an online retailer voucher equivalent to $30 as compensation. The study was approved by a research ethics committee at the authors' institution. Subjects were sat in a comfortable chair to minimize movement artifacts, approximately 47" (120cm) away from a 24" flat monitor that displayed the virtual agent. The room was quiet (but not soundproof) and dimly lit. The fNIRS probe positioned on their forehead was also covered with non-transparent fabric to prevent ambient light from reaching the sensors. Subjects were instructed to refrain from moving their limbs, frowning and talking while they carried out the experimental tasks.

To attenuate noise for post-hoc analyses (resulting from minor head movements, heartbeat and respiration), we applied sliding-window motion artifact rejection and raw data were low-pass filtered using a finite impulse response filter with order 20 and 0.1Hz cut-off frequency [4]. Additionally, we applied first order linear detrending on data from each channel [3]. HbO for each channel was calculated with re-

[2]Note that this approach to determining NF threshold for left-asymmetry is consistent with the original one of Rosenfeld et al. [27].

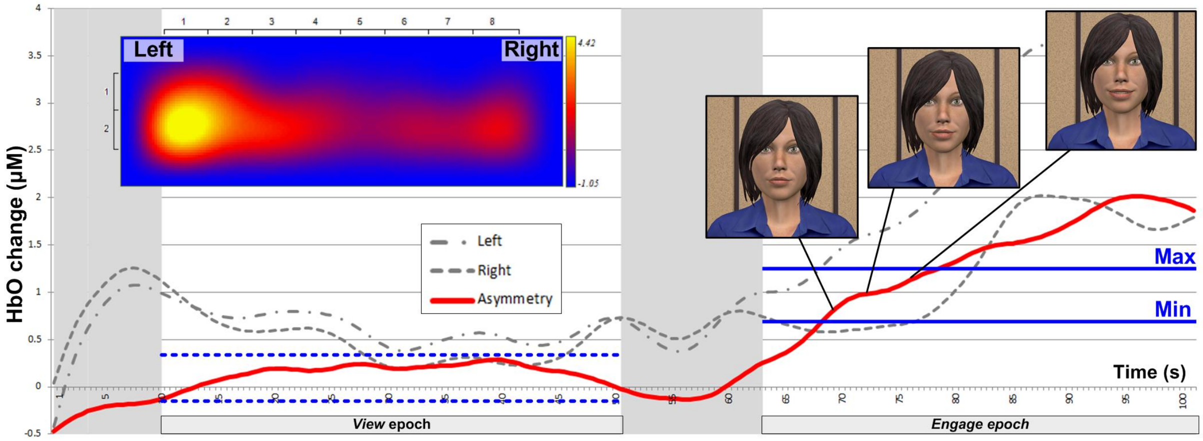

Figure 7: Example of a successful block where the asymmetry values in the *Engage* epoch are significantly greater than the ones in the *View* epoch, mapped onto the changes in the ECA's facial expressions. The topographic view (top-left) shows the 16 optodes' activation with a significant increase of left oxygenation during the *Engage* epoch.

spect to the baseline measured at the end of each rest period at the start of each data-collection block [4].

We extracted HbO values on each channel for *View* and *Engage* epochs within each block, using time synchronization markers. To accommodate for the approximately 7s delay in the hemodynamic response [8], we excluded the first 14 data points (corresponding to 7s sampled at 2Hz) of each epoch on each channel, and included 14 data points after the completion of the epoch (see Figure 10; also see [29] for a similar approach).

4. RESULTS

Out of the 80 blocks collected from the 10 subjects, 16 blocks were excluded due to minor malfunctions, excessive movement, and reported discontinuation of the tasks during the block. This left 64 valid blocks for analysis.

We defined block success as a statistically significant increase in average left-asymmetry during the *Engage* epoch compared to the *View* epoch within the same block[3]. To determine success, we used an independent t-test with a bootstrapping resampling method (1000 samples, 95% confidence intervals) on asymmetry scores in consecutive *View* and *Engage* epochs. Based on this criterion, 37 blocks were successful (58% of blocks).

On a subject basis, we judged a subject to have been successful in the NF task if at least 50% of their valid blocks were successful (consistent with previous literature [9]). According to this criterion, 70% of subjects achieved neurofeedback success. Median success ratio (i.e., the proportion of successful blocks per subject) was 50%. Each subject completed at least one block successfully, and one subject completed all of the blocks successfully.

[3]We elected to compare *View* and *Engage* epochs to determine block success, as opposed to simply comparing asymmetry during *Engage* against the baseline, because the visual stimulus during these epochs was very similar (and conceptually the same), and subjects' cognitive activity was controlled with explicit instructions (i.e., count during *View* and expressing positive thoughts during *Engage*).

The majority of excluded blocks were due to subjects abandoning the counting task due to its high subjective difficulty. Since *View* and *Engage* epochs were comparable in terms of the presented visual stimulus and adequately matched for task difficulty, the observed differences in the left-asymmetry signal within blocks can be attributed to switching from the neutral task to the control task during which the user's level of engagement with the ECA through positive thinking.

We calculated the effect-size measure r to characterize the substantive significance of increase in asymmetry scores between successive *View* and *Engage* epochs within successful blocks. The smallest r effect-size that could be detected was .34 (medium)[4], which can be interpreted as approximately 11.5% of variance explained in asymmetry values by the *Engage* task. This indicates that our experimental setup was sensitive in detecting increase in left-asymmetry during *Engage* epochs. The median r in successful blocks was .82

[4]We use Cohen's [10] conventions for interpreting the magnitude of the effect-size measure r.

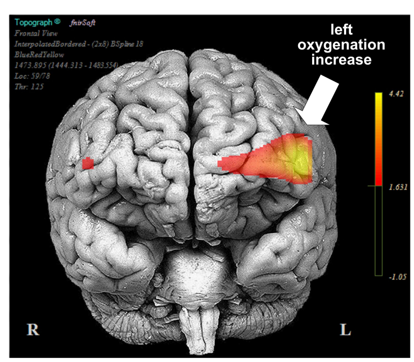

Figure 8: Topographic image of asymmetric dorsolateral prefrontal oxygenation increase during a successful NF epoch, overlaid on a brain surface image (generated using fNIRSoft by Biopac Systems).

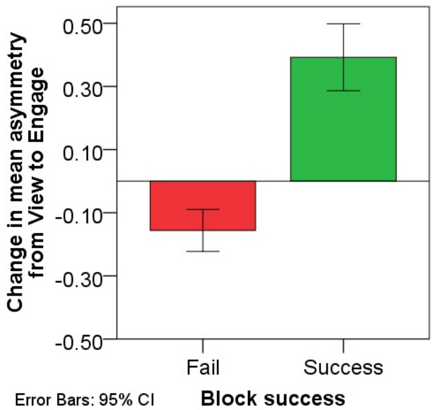

Figure 9: Change in mean asymmetry in successive *View* and *Engage* epochs, within successful and non-successful blocks. Note that non-successful blocks are characterized by a slight decrease in left-asymmetry during *Engage*, while successful blocks are characterized by a marked increase in left-asymmetry.

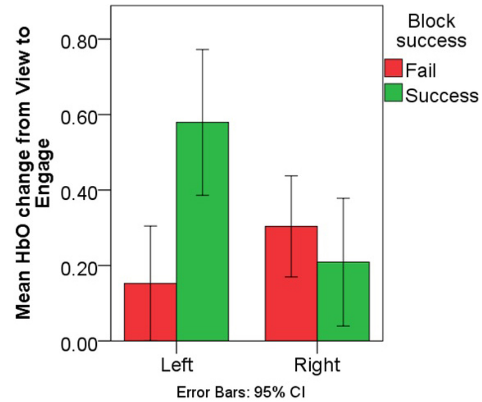

Figure 10: Differences in oxygenation increase from *View* to *Engage* epochs within successful and non-successful blocks, separately for the left and right sides. Note the marked increase in left-side oxygenation in successful blocks, while right-side oxygenation changes are at a similar level in the successful and non-successful blocks.

(large) with the interquartile range [.55; .87], which indicates that successful blocks were characterized with a marked increase in left-asymmetry during the *Engage* epoch, clearly distinguishable from the *View* epoch.

Figure 7 demonstrates how HbO changes on both sides contributed to the increase in asymmetry from *View* to *Engage* epochs. Left-side oxygenation was statistically significantly larger in successful blocks (M = 0.58, SD = 0.58) than in non-successful ones (M = 0.15, SD = 0.38), t(62) = 3.33, $p < .001$, r = .39 (medium). Conversely, right-side oxygenation did not differ significantly between non-successful blocks (M = 0.30, SD = 0.34) and successful ones (M = 0.21, SD = 0.51), t(62) = 0.40, p = .375, ns. These findings indicate that change in left-asymmetry in successful blocks was predominantly due to increase in left-side oxygenation during the *Engage* task, which is consistent with previous literature (see Figure 8).

Changes in left-asymmetry oxygenation between successive *View* and *Engage* epochs are presented in Figure 9, separately for successful and non-successful blocks. Non-successful blocks were characterized with a slight decrease in left-asymmetry during *Engage* (M = -0.16, SD = 0.17), while successful blocks were characterized with a marked and comparably large increase in left-asymmetry (M = 0.39, SD = 0.32).

Immediately after completing a block, subjects described their cognitive strategy to engage with the agent. Although the sample size of 37 successful blocks does not allow for a detailed content analysis, subjects' comments on their strategies suggest that thinking about personal positive memories was most conducive in capturing the ECA's attention, clearly illustrated on Figure 7 where as left-asymmetry increases during *Engage* epoch.

5. CONCLUSIONS

We have presented a novel experimental framework, which attempts to relate previous conceptualizations of the relationship between a user and an ECA to a single affective dimension. The expression of *approach* by users, captured through variations of their PFC asymmetry, could constitute an empirical basis to capture user disposition towards an ECA, in a more specific way than previously described concepts of engagement, empathy, alignment or rapport, which all refer to wider frameworks. What we have shown specifically is that, while approach was measured on the *user* side, on the *ECA* side it could be used to generate appropriate responsive behavior through the coordination of expressive parameters such as AUs and BAPs. These findings indicate that users were able to express a mental disposition to engage through the use of positive thoughts, which resulted in an appropriate response from the ECA congruent to their affective state. Although these early results are encouraging, further work should be dedicated to exploring the closed-loop aspects of these experiments, in particular the effects of ECA backchannel on users' attitudes.

6. REFERENCES

[1] I. Albrecht, M. Schröder, J. Haber, and H. Seidel. Mixed feelings: expression of non-basic emotions in a muscle-based talking head. *Special issue of Journal of Virtual Reality on Language, Speech & Gesture*, 8(4):201–212, 2005.

[2] R. Amini and C. Lisetti. HapFACS: An open source API/software to generate FACS-based expressions for ECAs animation and for corpus generation. In *Affective Computing and Intelligent Interaction (ACII)*, pages 270–275, 2013.

[3] H. Ayaz. *Functional near infrared spectroscopy based brain computer interface*. PhD thesis, School of Biomedical Engineering Science, 2010.

[4] H. Ayaz, M. Izzetoglu, P. A. Shewokis, and B. Onaral. Sliding-window motion artifact rejection for functional near-infrared spectroscopy. In *Engineering in Medicine and Biology Society (EMBC)*, pages 6567–6570, 2010.

[5] R. Beale and C. Creed. Affective interaction: How emotional agents affect users. *International Journal of Human-Computer Studies*, 67(9):755–776, 2009.

[6] E. Bevacqua, K. Prepin, R. Niewiadomski, E. de Sevin, and C. Pelachaud. Greta: Towards an interactive conversational virtual companion. *Artificial Companions in Society: perspectives on the Present and Future*, pages 143–156, 2010.

[7] T. D. Bui, D. Heylen, M. Poel, and A. Nijholt. Generation of facial expressions from emotion using a fuzzy rule based system. In D. C. M. Stumptner and M. Brooks, editors, *Australian Joint Conference on Artificial Intelligence*, pages 83–94, 2001.

[8] S. C. Bunce, M. Izzetoglu, K. Izzetoglu, B. Onaral, and K. Pourrezaei. Functional near-infrared spectroscopy. *Engineering in Medicine and Biology Society (EMBC)*, 25(4):54–62, 2006.

[9] M. Cavazza, G. Aranyi, F. Charles, J. Porteous, S. Gilroy, G. Jackont, I. Klovatch, G. Raz, N. J. Keynan, A. Cohen, et al. Frontal alpha asymmetry neurofeedback for brain-computer interfaces. In *Brain-Computer Interface Conference*, 2014.

[10] J. Cohen. *Statistical power analysis for the behavioral sciences.* Lawrence Erlbaum Associates, 1988.

[11] M. Courgeon, C. Clavel, and J. Martin. Modeling facial signs of appraisal during interaction: impact on users' perception and behavior. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 765–772, 2014.

[12] R. J. Davidson, P. Ekman, C. D. Saron, J. A. Senulis, and W. V. Friesen. Approach-withdrawal and cerebral asymmetry: Emotional expression and brain physiology: I. *Journal of Personality and Social Psychology*, 58(2):330–341, 1990.

[13] H. Doi, S. Nishitani, and K. Shinohara. NIRS as a tool for assaying emotional function in the prefrontal cortex. *Frontiers in Human Neuroscience*, 7, 2013.

[14] P. Ekman, W. V. Friesen, and J. C. Hager. *Facial Action Coding System: A Technique for the Measurement of Facial Movement.* Consulting Psychologists Press, Palo Alto, CA, 2002.

[15] S. A. Etemad, A. Arya, A. Parush, and S. DiPaola. Perceptual validity in animation of human motion. *Computer Animation and Virtual Worlds*, 2015.

[16] S. W. Gilroy, J. Porteous, F. Charles, M. Cavazza, E. Soreq, G. Raz, L. Ikar, A. Or-Borichov, U. Ben-Arie, I. Klovatch, et al. A brain-computer interface to a plan-based narrative. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1997–2005, 2013.

[17] K. Grammer and E. Oberzaucher. The reconstruction of facial expressions in embodied systems. *ZiF: Mitteilungen*, 2, 2006.

[18] L. Huang, L.-P. Morency, and J. Gratch. Virtual rapport 2.0. In *Intelligent Virtual Agents (IVA)*, pages 68–79, 2011.

[19] E. Krumhuber, A. S. Manstead, D. Cosker, D. Marshall, and P. L. Rosin. Effects of dynamic attributes of smiles in human and synthetic faces: A simulated job interview setting. *Journal of Nonverbal Behavior*, 33(1):1–15, 2009.

[20] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133, 1943.

[21] S. Moghimi, A. Kushki, S. Power, A. M. Guerguerian, and T. Chau. Automatic detection of a prefrontal cortical response to emotionally rated music using multi-channel near-infrared spectroscopy. *Journal of Neural Engineering*, 9(2), 2012.

[22] R. Niewiadomski, S. J. Hyniewska, and C. Pelachaud. Constraint-based model for synthesis of multimodal sequential expressions of emotions. *IEEE Transactions on Affective Computing*, 2(3):134–146, 2011.

[23] R. Niewiadomski and C. Pelachaud. Model of facial expressions management for an embodied conversational agent. In *Affective Computing and Intelligent Interaction (ACII)*, pages 12–23. 2007.

[24] M. Ochs and C. Pelachaud. Socially aware virtual characters: the social signal of smiles. *IEEE Signal Processing Magazine*, 30(2):128–132, 2013.

[25] M. Ochs, C. Pelachaud, and D. Sadek. An empathic virtual dialog agent to improve human-machine interaction. In *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 89–96, 2008.

[26] H. Prendinger and M. Ishizuka. The empathic companion: A character-based interface that addresses users' affective states. *International Journal of Applied Artificial Intelligence*, 19(3):297–285, 2005.

[27] J. P. Rosenfeld, G. Cha, T. Blair, and I. H. Gotlib. Operant (biofeedback) control of left-right frontal alpha power differences: Potential neurotherapy for affective disorders. *Biofeedback and Self-Regulation*, 20(3):241–258, 1995.

[28] A. C. Ruocco, A. H. Rodrigo, J. Lam, S. I. Di Domenico, B. Graves, and H. Ayaz. A problem-solving task specialized for functional neuroimaging: validation of the scarborough adaptation of the tower of london (s-tol) using near-infrared spectroscopy. *Frontiers in Human Neuroscience*, 8, 2014.

[29] P. Sarkheil, A. Zilverstand, N. Kilian-Hütten, F. Schneider, R. Goebel, and K. Mathiak. fMRI feedback enhances emotion regulation as evidenced by a reduced amygdala response. *Behavioural Brain Research*, 281:326–332, 2015.

[30] E. T. Solovey, A. Girouard, K. Chauncey, L. M. Hirshfield, A. Sassaroli, F. Zheng, S. Fantini, and R. J. Jacob. Using fNIRS brain sensing in realistic HCI settings: experiments and guidelines. In *the 22nd annual ACM symposium on User Interface Software and Technology (UIST)*, pages 157–166, 2009.

[31] K. Tai and T. Chau. Single-trial classification of NIRS signals during emotional induction tasks: towards a corporeal machine interface. *Journal of NeuroEngineering and Rehabilitation*, 6(1):1–14, 2009.

[32] R. Zhao, A. Papangelis, and J. Cassell. Towards a dyadic computational model of rapport management for human-virtual agent interaction. In *Intelligent Virtual Agents (IVA)*, pages 514–527, 2014.

[33] V. Zotev, R. Phillips, H. Yuan, M. Misaki, and J. Bodurka. Self-regulation of human brain activity using simultaneous real-time fMRI and EEG neurofeedback. *NeuroImage*, 85:985–995, 2014.