

Using near infrared spectroscopy to index temporal changes in affect in realistic human-robot interactions

Megan Strait¹ and Matthias Scheutz¹

¹Tufts University, 161 College Avenue, Medford MA 02155
{megan.strait, matthias.scheutz}@tufts.edu

Keywords: Brain Computer Interfaces, Functional Near Infrared Spectroscopy, Human Robot Interaction, Affect Detection, Signal Processing

Abstract: Recent work in HRI found that prefrontal hemodynamic activity correlated with participants' aversions to certain robots. Using a combination of brain-based objective measures and survey-based subjective measures, it was shown that increasing the presence (co-located vs. remote interaction) and human-likeness of the robot engaged greater neural activity in the prefrontal cortex and severely decreased preferences for future interactions. The results of this study suggest that brain-based measures may be able to capture participants' affective responses (aversion vs. affinity), and in a variety of interaction settings. However, the brain-based evidence of this work is limited to temporally-brief (6-second) post-interaction samples. Hence, it remains unknown whether such measures can capture affective responses over the course of the interactions (rather than post-hoc). Here we extend the previous analysis to look at changes in brain activity over the time course of more realistic human-robot interactions. In particular, we replicate the previous findings, and moreover find qualitative evidence suggesting the measurability of fluctuations in affect over the course of the full interactions.

1 INTRODUCTION

With recent advances in brain-imaging technology, inexpensive sensors are becoming increasingly accessible to researchers and consumers alike. Moreover, the production of sensors that are also small and wireless (in addition to being affordable) is promising for HRI, as that allows for the wearing of these sensors while performing a large set of activities without being intrusive. Examples of such devices include the Emotiv Epoch¹ and NeuroSky MindWave², which are two EEG headsets that measure electrical activity in the brain linked to states of excitement, attention, anxiety or cognitive load. Socially and affect-aware robots that can capture *and respond to* some of these states from a human have been found to be more effective in engaging people, e.g., (Szafrin and Mutlu, 2012). For these reasons, research on neurophysiological signals has been attracting the attention of researchers in the Human-Robot Interaction (HRI) community over recent years (Rosenthal von der Putten et al., 2013b).

Neural data, in particular, is extremely relevant

for HRI research in two main directions. First, it can complement traditional survey methods such as questionnaires and thus yield further understanding of users' genuine responses toward robots during the interaction, e.g., (Rosenthal von der Putten et al., 2013b; Strait et al., 2013a). Another potentially promising facet for neurological signal processing is affect detection in realtime, e.g., (Heger et al., 2013; Zander, 2009), so that the robot can react accordingly.

Within the HRI community, there is a small, but growing body of research employing brain-computer interfaces (BCIs) as a modality for both understanding and augmenting a person's experience (Canning and Scheutz, 2013; Zander, 2009; Zander and Kothe, 2011). Brain-based adaptivity of robotic agents has shown to yield performance and learning enhancements (Solovey et al., 2012; Szafrin and Mutlu, 2012). BCIs have also been used to further understand the user's perceptions of a robot, e.g., (Broadbent et al., 2013; Kawaguchi et al., 2012; Rosenthal von der Putten et al., 2013a; Strait et al., 2013a; Strait et al., 2014). In particular, a recent mixed-methods study employing a combination of brain-based and subjective measures reflected participants' affinity towards a robot (Strait et al., 2014; Strait and Scheutz, 2014).

¹<http://emotiv.com/epoc/features.php>

²<http://www.neurosky.com/Products/MindWave.aspx>

This study used a brain-imaging technique – functional near infrared spectroscopy (NIRS) – similar in basis to fMRI, but less restrictive (e.g., participants need not be confined to a tube).

However, the brain measurements were limited to brief (6-second), post-hoc re-exposure (viewing a series of images of the robots). Thus it remains unknown whether NIRS-based BCIs can capture changes in participants’ affect *over the course of the human-robot interactions*. Here we further probe the dataset collected in (Strait et al., 2014) to qualitatively evaluate changes in brain activity over interactions lasting two minutes in duration. While there are still a number of limitations of neurophysiology and NIRS that concern the usage of BCIs in realistic HRI interaction environments (Canning and Scheutz, 2013; Hoshi, 2011; Strait et al., 2013b), this investigation begins to bridge the gap between unrealistic and realistic interaction settings (e.g., viewing images of robots for several seconds versus unrestricted, live interactions).

2 RELATED WORK

Within the HRI community, there is a small, but growing body of research employing brain-computer interfaces as a modality for both understanding and augmenting a person’s experience (Canning and Scheutz, 2013; Zander, 2009; Zander and Kothe, 2011). Brain-based adaptivity of robotic agents based on students’ level of attention has been repeatedly shown to produce learning enhancements (Andujar et al., 2013; Szafir and Mutlu, 2012; Szafir and Mutlu, 2013). BCIs have also been used to further understand the user’s perceptions of a robot, e.g., (Broadbent et al., 2013; Kawaguchi et al., 2012; Rosenthal von der Putten et al., 2013a).

In particular, brain-based measures have been used for better understanding users’ perceptions of robots. For example, fMRI has been used to investigate emotional responses towards robots (Rosenthal von der Putten et al., 2013b) and towards a humanoid robot displaying affective gestures (Chaminade et al., 2010). Cooperation, rapport, and moral decision-making have been investigated using NIRS-based systems (Kawaguchi et al., 2012; Shibata, 2012; Strait et al., 2013a). A number of BCIs have been also been employed for affect detection, e.g., (Heger et al., 2013), by capitalizing on signal artifacts arising from facial expressions (Heger et al., 2011) or by targeting the ventromedial prefrontal cortex (Strait et al., 2013a).

While numerous exemplars of EEG-based BCI

systems exist for augmenting human-robot interactions, we focus here on NIRS in particular for the greater spatial resolution (which facilitates the targeting of ventromedial prefrontal cortical activity reflective of emotion regulatory processes). NIRS (functional Near Infrared Spectroscopy, also called ‘fNIRS’) is a neuroimaging technique similar to functional Magnetic Resonance Imaging (fMRI) that measures changes in blood flow corresponding to neural activity (Canning and Scheutz, 2013). From usability testing to animal-assisted therapy (e.g., (Kawaguchi et al., 2012; Shibata, 2012)), NIRS has served for over two decades as a quantitative metric in evaluating user workload, as an alternative interaction modality in assistive technologies, and more recently, as a passive input technique to adapt computer interfaces based on a user’s affective state (Zander and Kothe, 2011).

In comparison to other methods such as fMRI and EEG, NIRS-based systems have been reported as *better suited* for realistic settings, with the primary strength of being more robust to user movement (Cui et al., 2010a; Cui et al., 2011; Solovey et al., 2009). Recent work in HCI has further endorsed NIRS-based BCI as *suitable* for passive input to adaptive user interfaces based on improvements observed in *behavioral indices* of user performance (Solovey et al., 2011; Solovey et al., 2012). Moreover, although NIRS is vastly slower in temporal resolution compared to EEG, recent work has demonstrated reliable classification accuracies with under two-seconds delay from a response onset (Cui et al., 2010b). Thus it is potentially promising for use as input to robotic agents for adapting social behavior appropriate to user’s perceptions or affective state. However, few NIRS-based studies have investigated sequential or prolonged tasks (e.g., (Hoshi and Tamura, 1997)), rather NIRS is predominately employed in event-related experimental designs, limited to stimulus periods of sub-60s (Canning and Scheutz, 2013; Strait et al., 2013b; Hoshi, 2011). Thus it has yet to be shown whether NIRS is applicable to more realistic human-robot interaction settings.

3 MATERIALS AND METHODS

In a previous mixed-design human-robot interaction experiment, manipulating the type (MDS vs. PR2) of robot helper and modality of interaction (3rd-person remote vs. 1st-person remote vs. 1st-person collocated; see Figure 1), participants’ prefrontal cortical activity was recorded while they completed a set of two drawing tasks with each robot (Strait et al., 2014).



Figure 1: Subject perspectives for each of the three interaction modalities. Left – the 3rd-person perspective interaction (3R) condition with the PR2 as the robot helper. Right and center – the 1st-person perspective protocols, with the MDS robot helper remotely interacting (1R; center) and co-located with the participant (1C; right).

Manipulation of the robot helper was intended to vary the degree of the perceived human-likeness of robot (human-like versus mechanical). Thus the multi-dextrous social (MDS) robot was used for its relatively human-like appearance in contrast to the more typical appearance of Willow Garage’s PR2 robot. The three interaction conditions, on the other hand, were intended to modulate the directness (observatory versus participatory) and proximity (remote versus co-located) of the interaction.

To determine the effects of the aforementioned manipulations on participants’ perceptions of the two robot helpers (as indexed by prefrontal hemodynamic activity), Strait et al. additionally presented participants with a set of still images of each robot following completion of the four drawing tasks. This post-task exposure using still images allowed for the controlling of potential confounds including: artifacts from participant movement, unrelated hemodynamic changes (e.g., changes due to the drawing task), and unrelated physiological artifacts (e.g., low frequency signal drift).

A preliminary evaluation of the manipulation effects on the controlled post-exposure brain activity showed the MDS elicited significantly greater activity than the PR2 in the 1C condition. This effect was reversed in the 3R condition, which showed the PR2 elicited significantly greater activity. The effects in the first-person, remote (1R) interaction condition fell squarely between the two other interaction conditions, with no significant differences in response to either robot.

Here we extend the previous analysis, using the aforementioned NIRS dataset collected *during* the drawing tasks in (Strait et al., 2014) to investigate the changes in brain activity over the course of the full-length human-robot interactions.

3.1 Dataset

We utilized the NIRS dataset collected in the aforementioned IRB-approved study (Strait et al., 2014). In that study, 45 participants were instructed on four distinct tasks by two robot helpers (see Figure ??). Hardware and software issues of the NIRS equipment lead to failure to record or fully record seven participants’ interactions, resulting in a NIRS dataset of 38 participants. Subject demographics showed a 55/45 female-to-male ratio (21 female/17 male participants) and average age of 21.4 years ($SD = 4.1$).

3.1.1 Experimental Manipulations

Two primary manipulations were studied. To evaluate the effects of human-likeness, the Xitome Design’s Mobile Dexterous Social (MDS) robot and Willow Garage’s PR2 were used as the two robot instructors. They were chosen for their stereotypical robotic (PR2) and human-like (MDS) appearances and to avoid potential effects of height and girth of the helper (which are approximately equal between the two robots).

To measure the effects of interaction modality (comprised of participant perspective – first-person vs. third-person – and the robot’s presence – co-located vs. remote), three interaction conditions were created: (a) 3rd-person, remote (3R); (b) 1st-person, remote (1R); and (c) 1st-person, co-located (1C) – see Figure 1.

3.1.2 Data Acquisition

NIRS recordings of participants’ bilateral anterior prefrontal cortex were taken while participants interacted with each of the two robots using a two-channel NIRS oximeter. This placement of the NIRS sensors corresponds to areas linked, in particular, to emotion regulation (Chaminade et al., 2010; Ochsner et al.,

2012; Rosenthal von der Putten et al., 2013a; Strait et al., 2013a; Urry et al., 2006).

3.2 Signal Processing

Although the individual tasks were not fixed in duration, here we limit our analyses to the first two minutes of interaction for consistency for between-subjects comparisons. Thus, upon extracting the first two recorded minutes of each interaction, the NIRS dataset was then preprocessed in a similar manner as in (Strait et al., 2014): (1) conversion of raw light attenuation to changes in hemoglobin concentrations, (2) linear detrending to remove signal drift, (3) filtering of cardiac artifacts using a Savitzky-Golay low pass filter with degree 1 and cut-off frequency of 0.5Hz, and (4) correlation-based signal improvement ((Cui et al., 2010a)) to correct non-systemic artifacts (e.g., motion). Following, for each robot, we averaged across participants by condition, using one channel which corresponded to participants' oxygenated hemoglobin concentration changes in the left PFC. As there were two interaction tasks with each robot, this yielded one average timeseries per robot (2), per interaction modality (3), per task (2) for a total of 12 two-minute signals.

3.3 Statistical Inference

We first naively compared each pairing of signals (e.g., MDS vs. PR2 in 1C, first task) using repeated paired t-tests (Bonferroni corrected, $\alpha = .004$). As expected (given the signals were two minutes in duration), all pairings were statistically significantly different. As the tasks were free-form and not time-constrained, and moreover, as the placement of the NIRS probes were qualitative (aligned with the center of the forehead, atop the brow) – quantitative comparisons across participants would be confounded with differences in alignment of the interactions with the robots and alignment of the precise cortical area being sampled. Thus we proceeded with a qualitative discussion of their differences and the resulting implications for using NIRS to measure affective responses in more realistic settings.

4 DISCUSSION

In this investigation, we extended previous work to consider the following questions regarding the use of NIRS for evaluating human-robot interactions: (1) are there observable effects of interaction modality and human-likeness over more realistic task durations?,

(2) what, if any, are the effects of prolonged exposure or repeat interactions with a robot?, and (3) are these effects observable at the level of an individual? We first discuss the effects of interaction settings (modality and human-likeness of the robot agent), which mirror those observed in (Strait et al., 2014). We then examine the differences in activation between the first and second interactions performed with each robot, to discuss effects of repeated exposure. Lastly, we consider the findings from (1) and (2) from a within-subjects perspective.

4.1 Effect of Interaction Modality and Robot Human-Likeness

Consistent with the findings of (Strait et al., 2014; Strait and Scheutz, 2014), a correlation was observed between human-likeness and prefrontal hemodynamic activity (see Figure 2). Specifically, in the first-person, co-located interaction condition, the MDS robot elicited a significant *increase* in activity compared to the PR2. Whereas, in the third-person, remote interaction setting, the MDS elicited a significantly greater *decrease* in activity than the PR2. Furthermore, the first-person, remote settings showed comparatively no change (with relatively minor changes oscillating around zero). Interestingly, however, these results were only consistent with the previous when limited to the first 15-20 seconds of interaction.

In combination with the subjective responses reported in (Strait et al., 2014), the significant differences in neural activity to the two robots according to the interaction condition further underscores an effect of human-likeness and the corresponding perception of eeriness. In the 1PC condition, participants showed markedly greater activity in response to the MDS robot as compared to their response to the PR2. As the prefrontal cortex has been shown to be active in response to robots with high subjective ratings of eeriness (Strait and Scheutz, 2014), this suggests a strong emotional response was evoked in participants directly interacting with the very human-like MDS. Considering the subjective preferences reported in (Strait et al., 2014) also showed strong preferences for the PR2 versus the MDS (67% 1C participants versus 47% of 1R and 3R participants), this activity seems to be reflective of aversion to the more human-like robot when it is co-located with the participant. This effect seems initially contrary to the findings of (Broadbent et al., 2013; Lisetti, 2011) which suggest that a robot with a more human-like face is found to be more likeable. However, participants in Broadbent et al. and Lisetti interacted with computer-generated

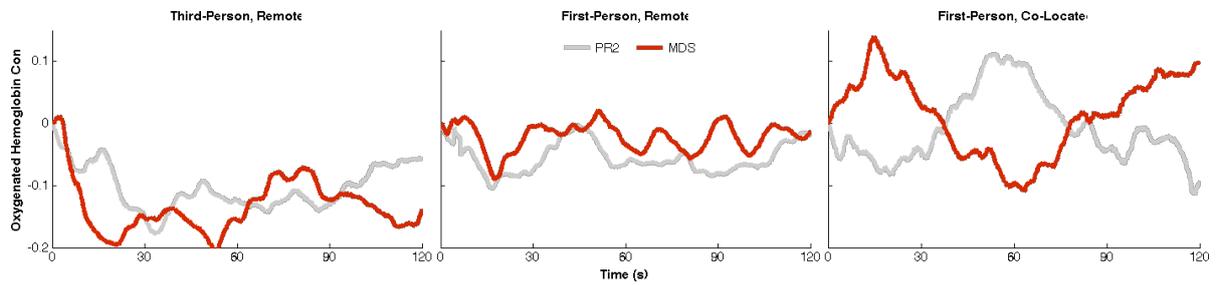


Figure 2: Average PFC activity during the *first* task with each robot (red depicts the interaction with the MDS, and gray, the PR2), by interaction condition. Left – the 3R interaction condition shows a decrease (initially greater in response to the MDS than PR2) in oxygenated hemoglobin across robots. Middle – similarly, the 1R interaction condition shows only slight changes over the two minutes of interaction for both robots. Right – the 1C condition shows an initial increase in HbO in response to the MDS and decrease in response to the PR2.

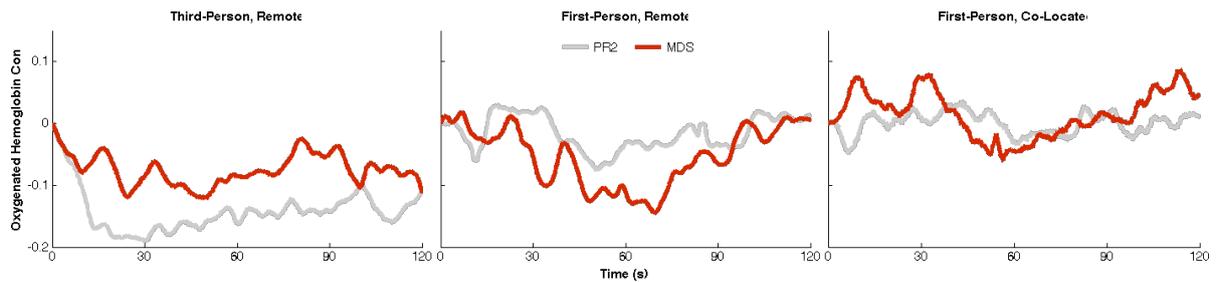


Figure 3: Average PFC activity over the course of the *second* interaction with each robot, by interaction condition. Left – the 3R interaction condition shows a decrease (greater in response to the PR2 than MDS) in oxy-hemoglobin across robots. Middle and right – the first-person interaction conditions show no major changes in hemoglobin in response to either robot.

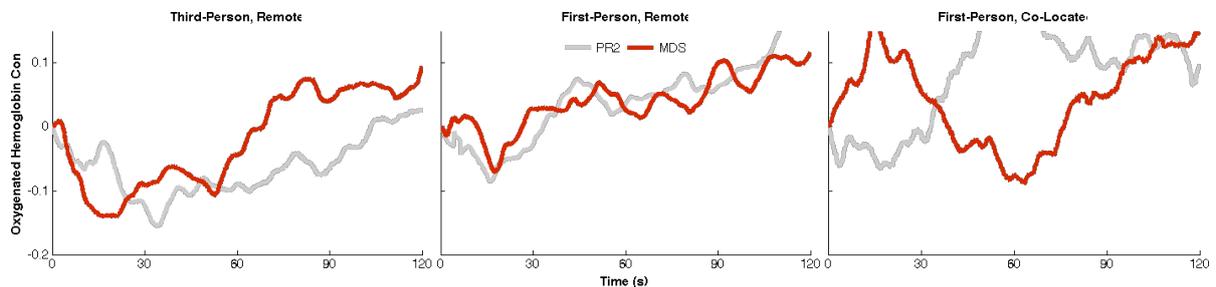


Figure 4: *Non-detrended* average PFC activity over the course of the first with each robot. Qualitatively the trends (decrease in hemoglobin in response to the MDS in 3R versus increase in hemoglobin in response to the MDS in 1C) are apparent.

avatars or a robot with a display screen face, rather than an embodied agent. Given similar settings (the 1R condition: interaction with the robots via Skype), participants seemed emotionally unaffected, perhaps due to relative reduction in robot presence.

Overall, these results replicate previous findings which suggest the correspondance of prefrontal activity with aversion (Strait et al., 2014; Strait and Scheutz, 2014) – in that it seems the eeriness of the MDS may elicit emotion regulation mechanisms in first-person interaction to reduce the unnerving effects of human-likeness and the corresponding eeriness. Whereas in a removed context such as that of observing video of the two – much like viewing a movie – the fear or anxiety elicited by the MDS’ eerie appearance may have been reduced or non-present. However, they also suggest that participants’ affective responses may change over the course of the interaction (e.g., the decrease in hemodynamic activity from 15s to 60s in the 1C interaction condition with the MDS; Figure 2, right). Due to this observation of significant change in the 1C condition, we next considered how the prefrontal activity during participants’ second interaction (i.e., second 2 minute task) with each robot compared to their first.

4.2 Effect of Repeat Interactions

The repeat interactions (the second task performed with each robot) or prolonged exposure to each robot showed significantly reduced responses (e.g., the magnitude of activity elicited in response to the MDS was decreased) in some settings and significantly enhanced in others (see Figure 3).

Figure 3 shows the prefrontal activity during the second interaction with each robot. In the third-person interaction condition, this repeat interaction shows the reverse trends from the first interaction: originally the MDS elicited greater negative change in comparison to the PR2 for 3R subjects. However, in the second interaction, participants show a more negative change in response to the PR2 (and the magnitude of the response to the MDS is reduced). In addition, participants in the 1R condition now showed a slight and slow decrease in HbO over the first 60+ seconds of the interaction. Activity elicited by the PR2 in the 1R condition still remained around zero. Participants in the 1C condition interestingly now showed a response similar to participants in the 1R condition. Here in the second task, we observe only a slight increase in HbO in response to the MDS followed by a slighter and slower decrease over the first 60 seconds. Whereas previously, the 1R condition showed relatively no response and the 1C condition showed a

severe increase in response to the MDS. Moreover, 1C participants showed a strong increase in HbO to the PR2 in the middle 60s of the first interaction, which is entirely absent from the second interaction.

While the changes occurring within the timeframe of the tasks (e.g., PR2 activity peak in HbO at 60s in the 1C condition) may be a function of the interaction (e.g., the PR2’s cooling fans suddenly turn on in the middle of the task), they nevertheless may be reflective of temporally-brief aversions of the participants to the robots. If the activity, for instance, in the 2nd task in the 1R condition is reflective of a growing aversion to the MDS or if the decrease in magnitude of brain activity in the 2nd task in the 1C condition is reflective of a decrease in aversion, this information becomes particularly relevant to how a robot might adjust it’s behavior. However, to deploy NIRS as a mechanism for adapting robot behavior, it is also necessary that these effects be observable at the individual level and not solely as an aggregate trend. Moreover, to influence a robot’s behavior in realtime, the delay in signal processing becomes an important consideration. Thus we next considered the persistence of the observed effects in two facets: (1) whether the effects are visible at the level of a single participant and (2) whether the effects are qualitatively observable in the absence of signal detrending.

4.3 Persistence of Effects

Since detrending, by definition, removes low-frequency signal drift, it necessarily requires a temporal delay on the order of the lowest-frequency signal artifact. While some work has shown using an exponential moving average with a 20s processing delay is sufficient to remove such artifacts (Cui et al., 2010a; Cui et al., 2011), this window still may exceed the duration of changes in affect. Hence, we first considered the effect of reducing the data preprocessing by eliminating signal detrending (see Figure 4). Qualitatively, the magnitude of the effects (e.g., severe increase in 1C response to MDS) are still observable. However, smaller effect sizes may be obscured by the low-frequency trends. For example, the 3R signal in response to the MDS is highly similar to that of the 1C condition from 60s to then end of the task. However, in the detrended signals, the two are much more dissimilar. Depending on the robot’s adaptive behaviors, misclassification of the 3R signal as indicative of aversion may or may not be important. Thus future work to develop computational methods for detection of such effects would need to address whether a change in signal was a result of signal drift or of a change in affect.

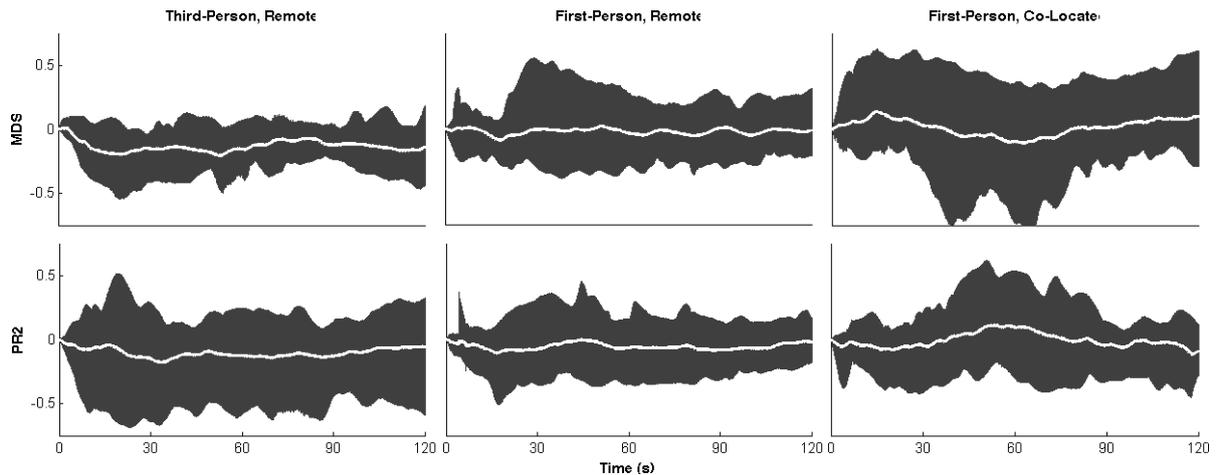


Figure 5: Range of hemodynamic activity in response to each robot across subjects. The average activity is shown in white, and the range (minimum HbO concentration to max. HbO concentration) for subjects in the given condition is shown in gray.

In addition, looking at the range (as opposed to global average) of brain activity shows high variation across participants and across conditions (see Figure 5). While the variation in hemodynamic activity across subjects qualitatively follows the trends captured by the global averages, a number of subjects (3 in each 1C and 3R conditions) in each interaction condition show relatively no significant changes in HbO throughout the task durations. One interpretation may be that such subjects did not have any affective response to the interactions. Further investigation with a larger population size or additional measures may help disentangle these responses. Conversely, there are a number of participants who show large signal spikes, suggestive that large motion artifacts were not adequately filtered. Perhaps the latter can be addressed through various approaches to motion filtering; however, state-of-the-art NIRS signal processing still suggests manual inspection or motion restrictions to adequately filter (Canning and Scheutz, 2013; Solovey et al., 2009; Strait et al., 2013b). While these analyses are qualitative in nature, they suggest that further investigation is necessary of the persistence of effects of robot appearance and interaction modality, and whether they are observable at a more micro level.

5 CONCLUSIONS

The aim of this study was to investigate the use of NIRS-based brain imaging for measuring human-robot interactions in more realistic settings than previously. Specifically, we qualitatively analyzed prefrontal cortical activity over the course of two-minute (semi) free-form interactions. Moreover, we evalu-

ated the changes in activity over repeat interactions and whether the effects observed were large enough to persist at an individual subject-level.

The findings are consistent with prior results, suggesting, in combination with subjective measures of preference, that PFC hemodynamics reflect a person’s aversion to a robot and are moderated by the settings of the interaction and the human-likeness of the robot interlocutor. But our results also suggest that participants’ responses fluctuate over the course of a task, and may diminish with prolonged exposure to or interaction with a given robot.

However, given the high variability and noise in the NIRS data, it is unclear whether these effects are reliably observable at an individual level. Thus the use of NIRS as a feasible realtime measurement for adapting robot behavior based on subject aversion may be too premature to attempt without more controlled investigations to better understand the individual variations in signal. Moreover, whether PFC activity is representative of a negatively-valenced affective response requires further investigation (either in combination with fMRI or with additional physiological measures) in order to disconfirm the presence of any confounds (e.g., signal artifacts or task-unrelated activity).

Despite current limitations to the use and deployment of NIRS in realistic, realtime settings; this evaluation provides an important bridge between high-controlled experiments showing observable effects over brief exposure (e.g., viewing images for several seconds) to actual, prolonged interactions. Future work will thus continue to investigate this relationship between signal changes and the extent to which variations of physical appearance of the robot have an influence on the perception of the interaction.

REFERENCES

- Andujar, M., Ekandem, J. I., Gilbert, J. E., and Morreale, P. (2013). Evaluating engagement physiologically and knowledge retention subjectively through two different learning techniques. In *HCI International*.
- Broadbent, E., Kuman, V., Li, X., Sollers, J., Stafford, R. Q., MacDonald, B. A., and Wegner, D. M. (2013). Robots with display screens: a robot with a more humanlike face display is perceived to have more mind and a better personality. *PLOS ONE*, 8(8):e72589.
- Canning, C. and Scheutz, M. (2013). Functional near-infrared spectroscopy in human-robot interaction. *Journal of Human-Robot Interaction*, 2(3):62–84.
- Chaminade, T., Zecca, M., Blakemore, S., Takanishi, A., Frith, C. D., Micera, S., Dario, P., Rizzolatti, G., Gallese, V., and Umiltà, M. A. (2010). Brain response to a humanoid robot in areas implicated in the perception of human emotional gestures. *PLoS one*, 5(7).
- Cui, X., Bray, S., Bryant, D. M., Glover, G. H., and Reiss, A. L. (2011). A quantitative comparison of NIRS and fMRI across multiple cognitive tasks. *NeuroImage*, 54(4):2808–2821.
- Cui, X., Bray, S., and Reiss, A. L. (2010a). Functional near infrared spectroscopy (NIRS) signal improvement based on negative correlation between oxygenated and deoxygenated hemoglobin dynamics. *NeuroImage*, 49(4):3039–3046.
- Cui, X., Bray, S., and Reiss, A. L. (2010b). Speeded near infrared spectroscopy (NIRS) response detection. *PLoS ONE*, 5(11).
- Heger, D., Mutter, R., Herff, C., Putze, F., and Schultz, T. (2013). Continuous recognition of affective states by functional near infrared spectroscopy signals. In *Proc. ACII*.
- Heger, D., Putze, F., and Schultz, T. (2011). Online recognition of facial actions for natural EEG-based BCI applications. In *Proc. ACII*, pages 436–446.
- Hoshi, Y. (2011). Functional near-infrared spectroscopy: current status and future prospects. *J. of Biomed. Op.*, 12(6).
- Hoshi, Y. and Tamura, M. (1997). Near-infrared optical detection of sequential brain activation in the prefrontal cortex during mental tasks. *NeuroImage*, 5(4):292–7.
- Kawaguchi, Y., Wada, K., Okamoto, M., Tsujii, T., Shibata, T., and Sakatani, K. (2012). Investigation of brain activity after interaction with seal robot measured by fNIRS. In *IEEE International Symposium on Robot and Human Interactive Communication*.
- Lisetti, C. L. (2011). Believable agents, engagement, and health interventions. In *Proceedings of the Human Computer Interaction International Conference*.
- Ochsner, K. N., Silvers, J. A., and Buhle, J. T. (2012). Functional imaging studies of emotion regulation: a synthetic review and evolving model of cognitive control of emotion. *An. of NY Acad. Sci.*, 1251:1–24.
- Rosenthal von der Putten, A. M., Kramer, N. C., Hoffmann, L., Sobieraj, S., and Eimler, S. C. (2013a). An experimental study on emotional reactions towards a robot. *International Journal of Social Robotics*, 5:17–34.
- Rosenthal von der Putten, A. M., Schulte, F. P., Eimler, S. C., Hoffmann, L., Sobieraj, S., Maderwald, S., Kramer, N. C., and Brand, M. (2013b). Neural correlates of empathy towards robots. In *Proc. HRI*, pages 215–216.
- Shibata, T. (2012). Therapeutic seal robot as biofeedback medical device: Qualitative and quantitative evaluations of robot therapy in dementia care. *Proceedings of the IEEE*, 100(8):2527–2538.
- Solovey, E. T., Girouard, A., Chauncey, K., Hirshfield, L. M., Sassaroli, A., Zheng, F., Fantini, S., and Jacob, R. J. K. (2009). Using fNIRS brain sensing in realistic HCI settings: experiments and guidelines. In *Proc. UIST*.
- Solovey, E. T., Lalooses, F., Chauncey, K., Weaver, D., Parasi, M., Scheutz, M., Sassaroli, A., Fantini, S., Schermerhorn, P., Girouard, A., and Jacob, R. J. K. (2011). Sensing cognitive multitasking for a brain-based adaptive user interface. In *Proc. CHI*.
- Solovey, E. T., Schermerhorn, P., Scheutz, M., Sassaroli, A., Fantini, S., and Jacob, R. J. K. (2012). Brainput: enhancing interactive systems with streaming fnirs brain input. In *Proc. CHI*.
- Strait, M., Briggs, G., and Scheutz, M. (2013a). Neural correlates of agency ascription and its role in emotional and non-emotional decision-making. In *Affective Computing and Intelligent Interaction*.
- Strait, M., Canning, C., and Scheutz, M. (2013b). Limitations of NIRS-based BCI for realistic applications in human-computer interaction. In *5th International BCI Meeting*, pages 6–7.
- Strait, M., Canning, C., and Scheutz, M. (2014). Investigating the effects of robot communication strategies in advice-giving situations based on robot appearance, interaction modality, and distance. In *Conference on Human-Robot Interaction (under review)*.
- Strait, M. and Scheutz, M. (2014). Measuring users' responses to humans, robots, and human-like robots with functional near infrared spectroscopy. In *Conference on Human Factors in Computing (under review)*.
- Szafir, D. and Mutlu, B. (2012). Pay attention!: designing adaptive agents that monitor and improve user engagement. In *Proceedings of CHI*, pages 11–20.
- Szafir, D. and Mutlu, B. (2013). Artful: Adaptive review technology for flipped learning. In *Proc. of CHI*.
- Urry, H. L., van Reekum, C. M., Johnstone, T., Kalin, N. H., Thurow, M. E., Schaefer, H. S., Jackson, C. A., Frye, C. J., Geischar, L. L., Alexander, A. L., and Davidson, R. J. (2006). Amygdala and ventromedial prefrontal cortex are inversely coupled during regulation of negative affect and predict the diurnal pattern of cortisol secretion among older adults. *J. Neurosci.*, 26(16):4415–4425.
- Zander, T. O. (2009). Detecting affective covert user states with passive brain-computer interfaces. In *Proceeding of Affective Computing and Intelligent Interaction*.
- Zander, T. O. and Kothe, C. (2011). Towards passive brain-computer interfaces: applying brain-computer interface technology to human-machine systems in general. *Journal of Neural Engineering*, 8(2).