Proceedings of Machine Learning Research 122:1–8, 2019 3rd International Workshop on Artificial Intelligence in Affective Computing

Sequential Dependence and Non-linearity in Affective Responses: a Skin Conductance Example

Jennifer Healey JEHEALEY@ADOBE.COM *Adobe Research*

Editor: William Hsu

Abstract

Individual affective responses frequently vary from the mean and often exhibit non-linear and time and sequence dependent properties. This paper examines the extent to which commonly made assumptions of linearity and sequential independence are valid using skin conductance responses to an acoustic stimulus as an example. We present 19 sessions of skin conductance traces where participants respond to five 50 millisecond acoustic bursts designed to elicit a startle. We show the data from the perspective of an online algorithm: individual responses, non-linear and dependent on prior events. We show that the coefficient of variation depends on sequence position and that these are large at 65%, 97%, 110%, and 100%. We discuss the risk of making inferences on single impressions.

Keywords: affect, sequential, independence, non-linear, response, startle, affect, skin conductance, variance, time-invariance, physiological

1. Introduction

Affective Computing takes an algorithmic perspective on data. Instead of looking for significant differences between populations, the goal often to infer affective state from a single individual. This is a challenging task for many reasons including acquiring data, labelling data and modelling people. One of the simplifing assumptions that is often made about people is that they act as linear, time-invariant (LTI) systems. This assumption is often made implicitly even though it is known not to be true in most cases. Even in Picard's seminal book on Affective ComputingPicard (1995), affective response is theoretically modeled as the additive response of a bell to repeated strikes. The bell model is that of a an LTI system where the response of is both consistent and proportional to the input. This paper strives to give examples of how people differ from this type of LTI model in one of the most basic types of responses, the startle response to an acoustic stimulus. In this experiment the acoustic burst is the analog to the hammer strike on the bell and the skin conductance is a readout of the "system function" of the participant.

2. Prior work

Skin conductance is a widely studied signal linked to the emotional aspect of arousalBoucsein (2012). It is well known that the skin conductance response is non-linear and is subject to habituation, a process that diminishes the amplitude of the response over timeThompson. It is also known that the skin conductance tends to recover from habituation which has led to a dual process theoryGroves (1970) where skin conductance is end result of both habituation and a sensitization (amplifying) process working silmultaneously. Recent research suggests that the mean amplitude results trends towards extinction until approximately the eleventh stimulation and then sensitization begins to dominate. We follow a protocol similar to Steiner and Barry and Steiner and Barry (2014) which use a series of acoustic sound bursts spaces at 5-7s intervals and 13 to 15 second inter-stimulus intervalsSteiner and Barry (2014) respectively. One of the limitations of this prior work research is that the presentation of the results consists of a report of the mean trend with associated standard error bars and does not show the degree of individual variation and the variety of the dynamic interaction of habituation and sensitization in individual traces. This paper focuses on showcasing individual variety with the goal of showing the importance of considering sequential context and of managing expectations around how reliably an algorithm might be able to differentiate an individual's emotional state based on limited observations. We present a number of different individual traces that exhibit early and extreme habituation, reactions to unstimulated responses and habituation and sensitization processes that do not follow the mean trend.

3. Skin Conductance Startle Response

Skin conductance is a metric that indirectly measures autonomic nervous system activity. Skin is normally an insulator and becomes conductive only when ionic (salty) sweat fills the sweat glands. When a person is startled, the central nervous system is activated and sweat suddenly fills the glands, causing the skin conductance to rise. An example of this is shown Figure 1 (a). The lower signal shows a time synchronized microphone sensor. The spike in the lower signal indicates the time of the acoustic burst. This time is marked with the first "x" on the upper skin response signal. The second "x" on the skin conductance signal shows the beginning of the rise, called the onset. The time between the stimulus and the onset is called the latency. Conductance continues to rise as the glands fill, eventually reaching a "peak" marked by a third "x." The vertical distance from the onset to peak is the amplitude. As the nervous system activation decreases, sweat is reabsorbed into the body and the skin conductance is diminished. This is the recovery, measured as time from the peak to the return to baseline (here shown as onset amplitude). Although there are many measurement variations, this one is commonBoucsein (2012).

4. Startle Elicitation Experiment

These traces were captured in a laboratory environment with each participant seated in a chair facing a blank wall. Participants were informed about the nature of the sensors and that they would be exposed to a series of sound bursts. Both skin conductance and sound were recorded using an Thought Technologies ProComp+ sensing system with the micrpohone attached to the system with a voltage isolater. This allowed the signals to be time synchronized. The Ag-AgCL button electrodes were placed on the underside of the middle section (middle phalanx) of the middle and index finger of each participants dominant hand. Participants were selected from a convenience sample of university staff and undergraduate and graduate students. Prior to beginning the experiment Participants were given approximately two minutes to sit and relax and become accustomed to wearing the sensors then we told them the experiment would begin and turned on the sensors and started a computer program that played a 50 millisecond burst of white noise at ten second intervals with a random delay of plus or minus 0 to 1 seconds added to the interval. The program generated five such bursts and then terminated at which point participants were informed that the experiment was over.

Figure 1: A comparison of an idealized train of responses that follow the LTI

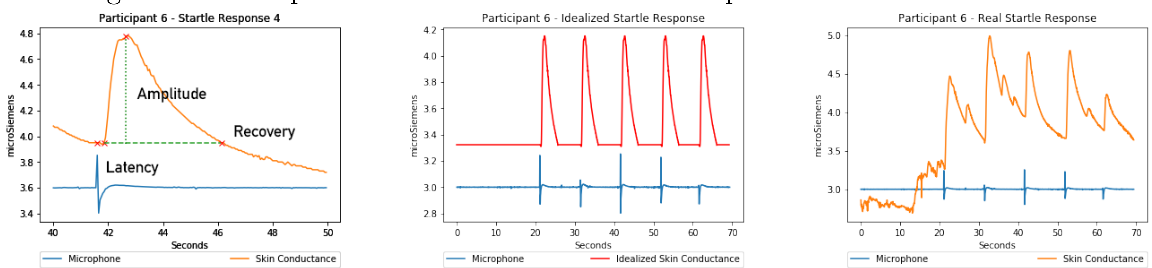

(a) Example Startle (b) LTI Model (c) Real Responses

Figure 2: Participant 12 (a) shows steady habituation while Participant 7 shows habituation followed by recovery.

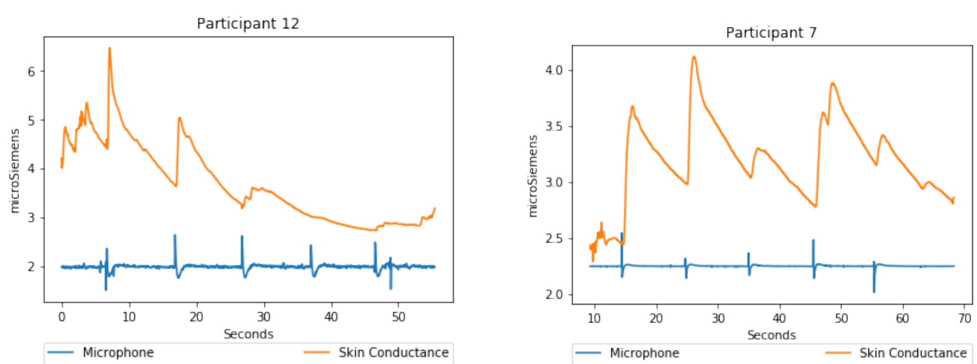

(a) Habituation (b) Habituation and Recovery

If the linear, time invariant assumption held for people, the response to each of these identical "impulses" should have been identical. Since all stimuli were played at the same volume and had the same duration, each response should be the identical and independent of its position int eh sequence. For Participant Six the response should look like an identical train of exact copies as illustrated in Figure 1 (b). This figure of the "Idealized Responses" was created by windowing the third response from Participant 6, from onset to recovery, and replicating it five times following each stimulus, resting on a line at the height of the onset of the response. This crafted response meets the assumptions of an LTI model. For contrast, the real response is shown in Figure 1 (c) where variation and other noise and unstimulated responses can also be seen.

It is well known that skin conductance orienting responses are not linear and as described earlier, scientists hypothesize that the diminishing effect, habituation and the amplifying effect, sensitization, work together to modulate the responses. Figure 2 (a) shows a consistent habituation response. Each successive startle response is diminished from the prior one. Figure 2 (b) shows potential habituation over the first three responses and then potential sensitization in the fourth response with continued habituation following. All sound bursts were identical (Section 5 deteails the rason for the different appearance of the microphone traces). It can be seen from these traces that habituation is not a consistent effect and that individual responses in some series do not show a strictly ordered attenuation.

Figure 3: Examples of traces with responses not directly attributable to the stimulus. Subfigure (b) shows possible habituation from the unstimulated response prior to the second stimulated response

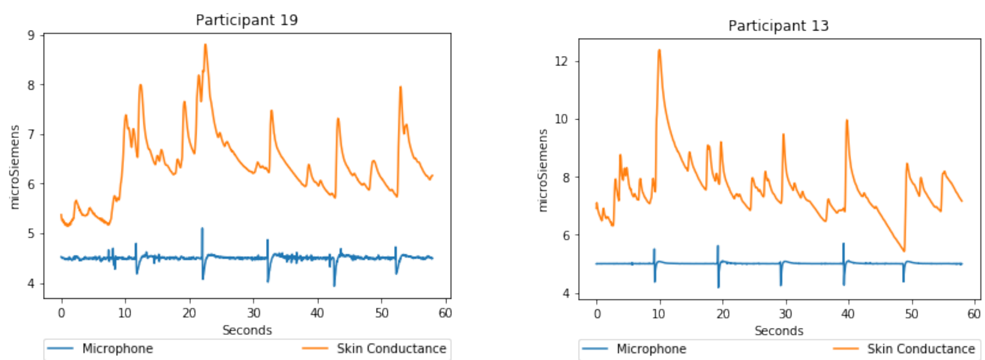

(a) Unstimulated Responses

(b) With Habituation

As further examples of the variation in real traces Figure 3 shows two sessions where unstimulated responses are abundant. These responses are difficult to separate and therefore difficult to quantify, especially without taking the history of the series into account. In Figure 3 (a) for example the second stimulated response onset occurs on the tail of multiple prior responses, none of which were intended by the experiment. In addition to multiple unstimulated responses, the trace from Participant 15 in Figure 3 also shows habituation and perhaps either sensitization or just variation of latter responses.

Habituation can also be rapid and extreme as shown in the trace from Participant 17 in Figure 4. The third, fourth and fifth stimuli seem to have no effect at all and the second response is only a small inflection. Without consideration of the high initial baseline and the substantial slow rising response to the first startle, the later responses might be interpreted as a person having no reaction to the stimuli. This lack of response might be misinterpreted as a psychopathic trait Lev if considered in isolation, but with the full history of the sequence a better hypothesis is perhaps that the person was initially very nervous about the experiment (thus the high initial baseline), but after the first noise burst decided it was really nothing to be alarmed about and went into a state of relief following exposure to the stimuli. The trace in Figure 4 (b) might at first not seem to exhibit the same degree

Figure 4: More extreme examples of habituation. Participant 17 shows no response after the initial response and Participant 25 shows only a minor response to the third stimulus and the fourth and fifth stimuli generate no response in the recovery of two prior unstimulated responses

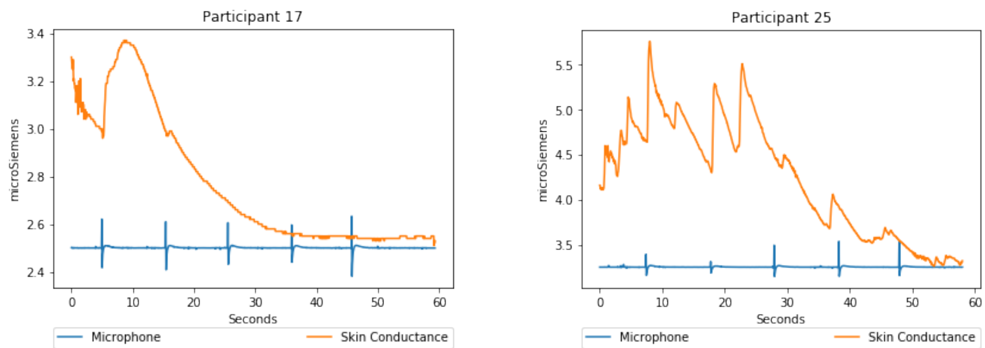

(a) Extreme Habituation

(b) Strong Habituation and Unstimulated Responses

of habituation as (a) however a closer look at the response to the third, fourth and fifth stimuli shows a very small response to stimulus three followed be a negative response to the fourth and fifth stimulus. This can be seen more clearly in Figure 5 (b) where the markers of the onset and peak are shown as detected by the signal processing. This trace actually shows a zero amplitude response to stimuli four and five and the recovery of prior unstimulated responses.

5. Signal Processing and Feature Extraction

We used a basic signal processing algorithm to extract the amplitudes of the responses. Both latency and recovery time are difficult to measure in the wild as often the timing of stimuli is not well known and recovery is often interrupted. We captured both skin conductance and acoustic signals at 20 samples per second (50 millisecond intervals), which was often slightly too low to capture the peak of the acoustic burst. We extracted the amplitude of the stimulated responses by first windowing the four seconds of data following the acoustic burst then finding the minimum and the maximum value of the windowed signal. These results were checked manually. In rare instances where the results were incorrect, due for example to a second unstimulated response occurring in the window (see Figure 3 (a) response or the peak being slower that anticipated as in Figure 4, the correct values were identified by hand. Two examples of the results of this process are shown in Figure 6(a) where the trace in (b) shows two negative results due to two zero amplitude responses on the recovery leg of prior responses. All such negative results were corrected to zero.

Figure 5: Examples of extracted amplitudes from two traces.

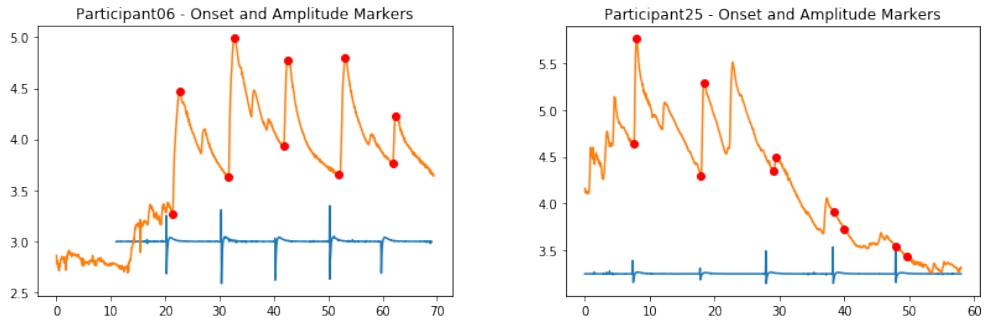

(a) Amplitude Extraction (b) Negative Amplitude Example

6. Results

The mean results of our experiment are presented in Figure 6. These values could be said to show a typical habituation response. The magnitude of the the first response is significantly larger than the subsequent responses with values of p=0.01, p=0.00, p=0.00, p=0.01 for each of the paired two-sided test for the null hypothesis for two related or repeated samples having identical average (expected) values for parings of the first response with the second, third, fourth and fifth response respectively. These results are in line with others that have been reported in the literature for similar studies Steiner and Barry Steiner and Barry (2014) Walker et al. with the exception of a higher degree of variation being seen for this study. This is potentially due to a less controlled lab setting, the small number of participants and the fact that participants were not screened or excluded based on caffeine intake, lack of sleep or other medications that might impact the results. This potentially makes the results of this experiment more contaminated, however it adds to the finding that these responses are likely sensitive to many factors and that one should expect that it is possible that not all relevant contextual variables are known for studies "in the wild."

7. Conclusions and Discussion

This paper gives a detailed presentation of an acoustic startle experiment from an affective computing perspective showing the variety of response across individuals and the ways that these signals are non-linear and sequential depdendent. The goal here is to show that results based on LTI assumptions are likely to make mistakes. Taken out of context, an algorithm might assume the "zero" responses from Participant 17 or 25 might indicate psychopathic tendencies rather than habituation. Additionally an algorithm might count the number of unstimulated responses in Participant 13 and 19 and decide that these people had higher anxiety levels than others or were less resilientWalker et al.. But we don't know this. It could well have been that they had just had coffee or that there was some other noise or distraction in the room at the time, contexts not prsent in more sterile lab studies. Affective processes, even those that regulate the as basic an impulse as acoustic startle response, are still not well modeled, even after over a hundred and fifty years.

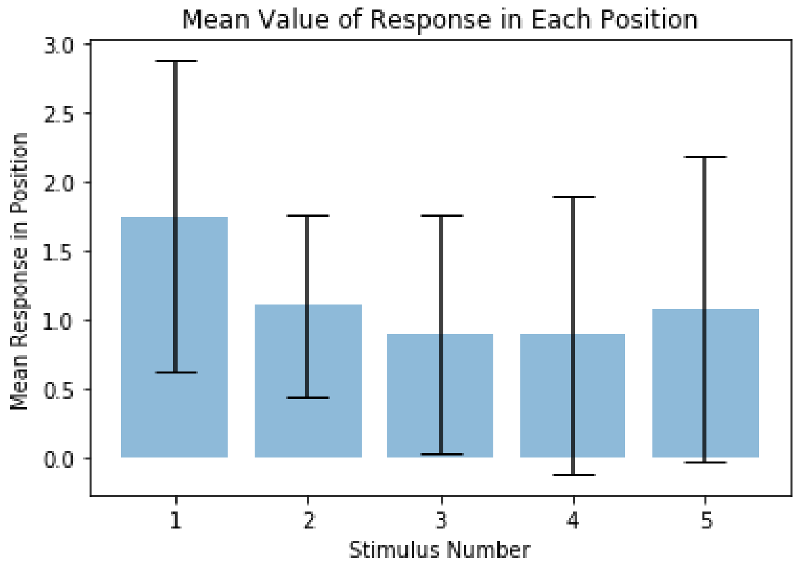

Figure 6: The Means and the standard deviations

It is naive to assume that affective computing is likely to be perfect in the near future given the variation of individual responses and the extent that they are impacted by prior and unrecorded events. Without much more sophisticated models that take into account a variety of contextual factors, we will fail.Barrett (2017). Current news articles make statements like: "software can scan a conversation between a woman and a child and determine if the woman is a mother, ... whether she is angry or frustrated or joyful. Still others can do so from facial expression."Dwork (2018)Khatchadourian (2015). While is is possible that an algorithm could recognize this situation correctly (e.g. the person was the mother and was perhaps speaking an a uniquely child-directed prosidic tonefer, it is certainly not the case that any algorithm can do this with 100% accuracy across individuals in the wild. The person could also easily be a nanny or other relative using "motherese" and conversely it is also highly likely that a mother-child relationship could exist if the computer did not detect it. Facial expression software is also not 100% accurate and often struggles to discern a smile from grimaceKhatchadourian (2015). And even if smiles were a bulletproof indicator of Joy, which they are not,Perepelkina and Astakhova (2018) even Ekman himself is moving away from the concept of "basic" emotions to emotion spectrums. Ekman (1992) Group (2016). It is very dangerous to assume that computers have near perfect emotion recognition for individuals, although still more articles imply that they somehow have "super-human" abilitiesCarrie (2017).

Emotion recognition in individuals from windowed online data is bound to be fraught with errors. In this paper, we wanted to demonstrate how much people can vary, even over short time windows and how much sequential context and prior events matter in interpreting signals. Looking at means across people can give lead to great scientific insight. Making recommendations (placing ads) with algorithms that simplistically infer a mean class can also result in significant revenue gains gain businesses on average. Means are useful, but they are not the whole story. In this paper we have shown examples of how the mean response is often very different from any individual response and we argue that making inferences about specific individuals based on single responses can easily lead to errors.

References

A cross-language study of prosodic modifications in mothers' and fathers' speech to preverbal infants. 16.

Lisa Feldman Barrett. The theory of constructed emotion: an active inference account of interoception and categorization. *Social cognitive and affective neuroscience*, 12(1):1–23, 2017.

Wolfram Boucsein. *Electrodermal Activity.* Springer Science and Business Media, 2012.

Carl Carrie. On affective computing: Past imperfect, future impactful. Hackermoon, November 2017.

Cynthia Dwork. Theory for society: Crypto on steroids. In *In Proceedings of the 38th International Cryptology Conference*, Santa Barbara, CA, USA, August 2018.

Paul Ekman. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200, 1992.

The Paul Ekman Group. Introducing the atlas of emotions. online, May 2016.

Richard F. Groves, Philip M. Thompson. Habituation: A dual-process theory. *Psychological Review*, 77(5):419–450, 1970.

Raffi Khatchadourian. We know how you feel. *The New Yorker*, January 2015.

Olga Perepelkina and Kristina Astakhova. 3 most widespread myths about emotions in AI. Data Driven Investor, October 2018.

Rosalind W. Picard. *Affective Computing.* MIT Press, Cambridge, Massachusetts, 1995.

Genevieve Steiner and Robert J Barry. Exploring the mechanism of dishabituation. *Neurobiology of Learning and Memory.*

Genevieve Steiner and Robert J Barry. The mechanism of dishabituation. *Frontiers in integrative neuroscience*, 8(14), Feb 2014.

Richard F Thompson. Habituation: a history. *Neurobiology of Learning and Memory.*

Frederick R. Walker, Ashley Thomson, Kane Pfingst, Elke Vlemincx, Eugene Aidman, and Eugene Nalivaiko. Habituation of the electrodermal response – a biological correlate of resilience? *PLoS ONE*, 14(1).