

PokerMetrics: Stress and Lie Detection through Non-Invasive Physiological Sensing

Michael Sung and Alex (Sandy) Pentland

MIT Media Laboratory, Human Dynamics Group, 20 Ames Street, Cambridge, MA

[\[msung@sandy\]@media.mit.edu](mailto:msung@sandy@media.mit.edu)

Abstract

In this study, we demonstrate that we can use non-invasive physiology sensing to detect stress and lying, within the context of poker settings. We show how simply derived non-invasive physiological features such as voice pitch variation, skin conductance peaks, and heart rate variability are correlated to a number of high stress situations found in real money no-limit Hold'em tournaments. Using these features, we can develop simple linear models that can be used to identify stress and bluffing.

1. Introduction

The National Institute of Occupational Safety and Health states that stress is becoming the most prevalent reason for worker disability. A 1992 UN report called job stress "The 20th Century Epidemic", while the World Health Organization stated in 1996 that stress was a "World Wide Epidemic". Researchers estimate job stress costs American industries between \$200 and 300 billion annually. Given the mounting social costs of stress, the possibility of automatically identifying and monitoring stress levels for intervention purposes is compelling. This is particularly true for people who regularly work in high stress environments, such as financial traders or emergency workers. In such situations, faulty performance as a result of acute stress can lead to million dollar losses of or even the loss of life.

Given the implications of stress on work performance, there has been recent interest in monitoring the work performance of individuals under stress. Specifically, studies in behavioral finance present psychophysiological evidence that even the most seasoned securities trader exhibits significant emotional response as measured by elevated levels of skin conductance and cardiovascular variables during certain transient market events [Lo & Repin 2002, Lo et al. 2005]. Other studies have corroborated the evidence linking emotion with trading performance [Steenbarger 2002]. These studies indicate that psychophysiology and stress are intimately linked, and it is possible to infer one from the other.

As a baseline study to determine if it is possible to correlate non-invasive physiology with a variety of important contextual measures such as stress, attention, and interest, a pilot study has been initiated to establish the link between stress and physiology. This study was started to look at player stress within real live poker game scenarios, namely no-limit Texas Hold'em tournaments. In such tournaments, one can lose an entire bankroll in a single hand, termed "going all-in". This has great potential in creating live stressful situations with objective outcomes in a highly controlled setting that can be monitored easily.

We wish to demonstrate that non-invasively derived physiology and behavioral sensing (motion, skin conductance, heart rate, and temperature/heat-flux) can be correlated with various stressful events in poker tournaments. Specifically, we will be looking at subjective reports of stress levels, bluffing, and all-in situa-

tions as outcomes and correlating these situations with physiology features aggregated over the hand.

2. Background

Whereas traditional clinical physiology monitoring focuses on identifying accurate physiological responses (e.g. ECG traces during a heart arrhythmia), long-term monitoring enabled by minimally invasive sensing provides the ability to correlate contextual measures over time to a person's behavior and internal state. Research has shown that it is possible to correlate minimally invasive physiology measures to identify notions of a person's intentions and affective state such as interest, happiness, and stress [Picard 2001, Picard et al. 2001].

In particular, accurately identifying human intention such as lying has been particularly contentious given its inherently subjective nature. The concept of the 'lie detector' has always captured the imagination and interest of the popular press since it was invented over a hundred years ago in 1902 by James Mackenzie. The prototypical version invented by Mackenzie, was also called the polygraph tool because it looked at a range of physiological phenomena such as a person's heart rate, breathing rate, blood pressure, and skin conductance while a person is questioned. Though a report released in 2002 by America's National Academy of Sciences indicated that polygraphs are not completely reliable (better as a measure of stress than veracity), there is no doubt that a variety of physiological phenomena can be correlated to lying.

If such an ambiguous thing as telling the truth can be predicted and correlated to basic physiological features such as those analyzed by a polygraph, it is not a large leap of faith to consider using physiology to quantify a person's other internal states such as stress, frustration, interest level, excitement, attention, drowsiness, and affect [Picard 2001, Picard et al. 2001]. In fact, there is research to show that these things can in fact be accurately quantified, and furthermore that these measures can demonstrate a high degree of correlation and concordance among interacting groups of people, whether being stimulated while passively watching a movie [Madan et al. 2004] or the highly involved interaction dynamics of a psychiatric patient and a therapist [Marci 2002].

3. Experimental Methodology

PokerMetrics Subject Selection and Protocol

The study population was drawn from the MIT community who answered recruitment emails sent to the MIT poker mailing list. MIT COUHES approval and written consent was obtained from all participants. All of the subjects were college-age, healthy adult male individuals drawn from the undergraduate and graduate schools at MIT. All subjects were self-purported poker aficionados who play semi-regularly, though the skill level ranged from beginners to players who play for profit in high-stakes games.

For the poker physiometrics study, the physiology and behavioral state of players were taken over 31 tournaments, representing 401 hands of poker. Every subject was paired up with another player to play real live-money games of no-limit Texas Hold'em in heads-up style tournaments. While the players were engaged in the tournament, which typically ranged in duration between 10-30 minutes, the physiology and behavioral state of each player was recorded and annotated using a LiveNet system, a flexible sensing platform [Sung 2005].

As financial interest was instrumental in generating the stressful situations for this study, the players were asked to play with real money buy-ins in winner-takes-all tournaments. Subjects were paired up to other players with similar financial risk thresholds (with a buy-in between \$5 to \$20/tournament). The player's self-reported skill level and risk tolerance were also recorded in order to match players as well as to provide additional baseline information.



Figure 1: Subject wearing the non-invasive sensors (heart rate chest strap, finger skin conductance electrodes, audio microphone, and integrated motion/heat flux/skin conductance monitor) along with the LiveNet platform running BioRecord during a session of the PokerMetrics Study. We use the collected physiology to correlate to manually recorded subjective outcomes such as stress level, interest, and bluffing.

The BioRecord application was used to record multimodal physiology through the duration of each hand, as well as to timestamp interesting annotations during the match. This data included heart rate, skin conductance, temperature, heat flux, speech features, and movement data. The recordings were done on a per hand basis per player. The multimodal data streams of each player were also synchronized to each other.

PokerMetrics Subjective and Objective Outcome Data

In addition to the physiological and behavioral data collected, each subject was asked to fill out a tournament diary at the end of each hand. This form asked information regarding both subjective ratings of the player's mental state as well as objective information from the hand. The subjective measures included a player's self-reported interest, excitement, happiness, and stress levels, all rated on a scale of 1-10 (the subjects were asked to rate their current state relative to the typical range of emotional

activity that they perceive playing poker). Objective outcome data on whether the individual won or lost, whether they bluffed (i.e. small bluff, medium bluff, large bluff, or semi-bluff), whether there was an all-in hand initiated, whether the hand was a bad-beat (defined as when a heavily mathematically favored hand results in a loss), and the amount of money wagered were also recorded. This combination of subjective measures as well as objective outcomes were correlated with the physiology collected. For the analysis, we focused on the situations in the tournament which are traditionally considered high stress (large bluffing, all-in, bad-beat, and subjective stress situations with a rating of greater than 7 on a scale of 1-10).

PokerMetrics Human Subjects Approval

To obtain human subjects approval, we needed to comply with both MIT's Committee on the Use of Humans as Experimental Subjects (COUHES) Department. For this, we needed to provide documentation detailing the experimental protocol, subject consent procedures, and subject recruitment. It was necessary to explicitly describe every type of data collected from the subjects as well the equipment used in the study. Also of particular interest was to make sure the protocol passed the Health Information Privacy Act (HIPA) that protects an individual's right to maintain their health information privacy. In order to comply with these regulations, we had to make sure that personally identifiable subject information was not included in the data that was collected.

4. Data Analysis

The physiological and contextual data for each hand of a particular tournament was individually segmented and synchronized. For each hand, a variety of features derived from the heart rate, skin conductance, temperature, audio, and accelerometer data was calculated over the duration of the segment, which typically lasted between 30 seconds for an uneventful hand up to a few minutes for an involved hand.

Given the short duration of each hand, the window for these feature calculations was the entire hand length. Standard heart rate measures such as the average rate, standard deviation in RR intervals, min/max ratios, and frequency-domain energy and entropy ratios were calculated from the hands. For skin conductance, heat flux and temperature, average tonic level, slope variation, and threshold peaks counts were calculated. For the audio features, energy, fundamental frequency, spectral entropy of voiced segments, and variation in these measures over time were calculated, as well as speaking dynamic features such as fractional speaking time, and voicing rates. Motion features such as energy, max/min ratios, entropy, and spectrographic features were calculated.

For each individual's hand, the physiology and behavioral features were computed over the entire hand during the course of play. Two typical hands are shown in the figure below.

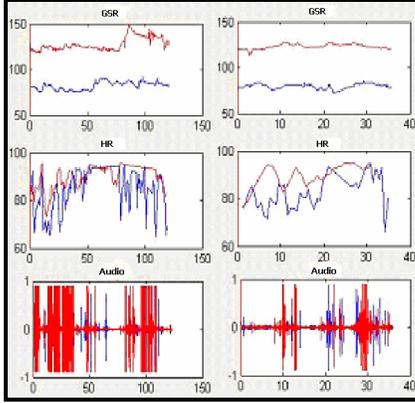


Figure 2: Skin conductance (top), heart rate (middle), and audio (bottom) data for a stressful all-in hand (left) vs. a non-eventful hand (right). As can be seen, the all-in play midway through the stressful hand was anticipated by the blue player who initiates it, shown as a small ramp in the skin conductance trace in the top left graph. This causes a large spike in the skin conductance of the red player a little bit later. We can also see that in general there was increased skin conductance, heart rate variability, and voice activity in the stressful hand relative to the non-stressful hand.

The hand on the right represents a non-eventful hand synchronized between two players in a tournament. We can see that the skin conductance did not change very much, and there were intermittent spurts of talking. In contrast is the last hand of the tournament between the two same players. The blue player first starts to become stressed (as indicated by the small rise in skin conductivity midway through the hand). A little bit later, he goes all-in, prompting a very large ramp in skin conductance by the red player while he ponders his options. Also note the very large differences in heart rate variability between the two hands (there was much more variability in the stressful hand). Finally, both the magnitude and amount of talking in the stressful hand was noticeably larger in the non-stressful hand.

The features were extracted by partitioning the data of a particular hand into windows (specifically before the resolution of the hand as well as after hand). This was because the voicing and motion dynamics were very different in these regimes since people tend to be more quiet and motionless in stressful situation before hand resolution, and exactly the opposite following the end of the hand.

All calculated features over each hand were normalized and z-scored to find the correlations to the outcomes. There were strong physiological correlations to most of the high-stress situations during play, most notably for all-in plays as well as for large bluffs. For large bluffs in particular, it was found that frequency-based heart variability features such as the LF/HF ratio, fraction speaking time, and motion energy were the most highly correlated features, as shown in Table 1 below:

Bluffing Top Features (n=32)	Statistics	
	R	p
1) HRV (LF/HF ratio)	0.72	0.018
2) Fraction Speaking Time	-0.68	0.021
3) Motion Energy	-0.65	0.03

Table 1: Top features correlated with bluffing. The heart rate variability LF/HF ratio, fraction of speaking time, and motion energy were the features most highly correlated to bluffing.

The LF/HF ratio PSD feature was often interpreted as a measure of the relative sympathetic to parasympathetic activity of the autonomic nervous system. The LF power was generated mainly from sympathetic activity, with baroreceptor modulation being a major component in LF power. The HF power, in contrast, was derived from vagal, or parasympathetic activity which is modulated by respiration. Thus, LF/HF ration represents a good indicator of sympathetic-vagal balance, and is used to assess the balance of the autonomic nervous system. As discussed in Section 4.5, studies have shown that mental stress increases LF activity and decreases HF activity. It seems that bluffing, insofar as it induces a stressful situation on a subject, results in a significant increase in LF power, and a slight decrease in HF power. Because of the increase in LF power, the LF/HF ratio was seen to increase during bluffing and other stressful situations.

The fraction of speaking time feature was calculated over the window of the recorded hand with 15 seconds from the end removed (which contains changes in physiology and behavior following the resolution of the hand). This feature was negatively correlated to bluffing.

The third feature, motion energy, was the mean energy calculated over the hand frame and was also negatively correlated to bluffing. We find that there was a significant reduction in the mean motion energy during a hand where the player was bluffing.

The later two features (speaking time and motion energy), are anecdotally well known to be associated with bluffing in the poker community [Caro 2003]. Specifically, when a person was bluffing, he or she consciously tries to reduce all interaction, including to stop speaking and to become very still, so as not to give a reason for the opponent to call the bluff. It is reassuring that this intuition holds true quantitatively in the features that were correlated well with bluffing.

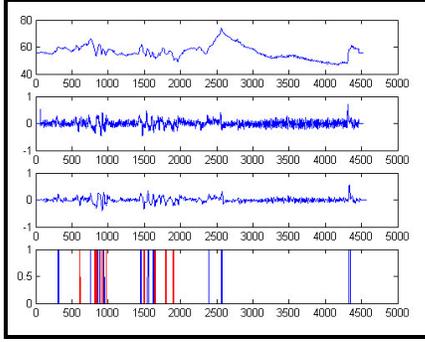


Figure 3: Raw skin conductance signal trace from the PokerMetrics study showing a response to a number of stressful situations (top), normalized difference (mid top), filtered difference (mid bottom) and negative (red) and positive (blue) thresholds (bottom). The thresholds can be set at different trigger levels to provide a gauge to the number of various-intensity stress or emotion response peaks that occur over a period of time.

All-in situations, where a player was faced with, or initiated a play where the rest of their money was in jeopardy, was another stress situation. For all-in hands, skin conductance peaks were the most highly correlated features. Below is a figure showing a hand where an all-in play was initiated a little after half-way through the hand. We see that there was a large ramp in skin conductance due to sympathetic arousal from the surprising and stressful situation that has transpired. By taking a running difference measure of the normalized skin conductance (scaled by the maximum range of the data), we get an approximation for the instantaneous change in slope of the original signal. This signal can then be smoothed by applying a low-pass filter, and then thresholded at different levels to get a sense of the number of peaks and valleys in the original skin conductance signal. Small thresholds catch smaller peaks, and larger thresholds can catch a very large skin conductance peak which turns out to be well correlated to high stress events.

The most basic HRV measure of the standard deviation of RR intervals was the second most correlated features. Voice energy also ended up being the most highly correlated voice feature.

All-Ins Top Features (n=39)	Statistics	
	R	p
1) Skin Conductance Peaks	0.75	0.013
2) Std. Dev. RR Intervals	0.71	0.017
3) Voice Pitch Variation	0.67	0.025

Table 2: Top features correlated with All-In plays. The number of skin conductance peaks, the std. dev. of the RR intervals of the heart, and voice pitch variation as measured by the std. deviation of the formant frequency were the most correlated features.

The bad beats results were less correlated to physiology and behavior during the hand, but very correlated to behavior after the hand has transpired. All bad beats had voice and motion energy as well as the standard deviation in the location of the largest auto correlation peak (essentially pitch variation). This makes sense; while the hands varied widely, the defining characteristic of a player being beat was to be very vocal (and whiney about it), as well as to move around in an agitated way following a hand. Statistical significance was an issue since there were relatively few recorded bad beat plays (around 19 instances in 401 hands of poker).

Bad Beats Top Features (n=19)	Statistics	
	R	p
1) Voice Pitch Energy	0.70	0.04
2) Motion Energy	0.63	0.06
3) Voice Pitch Variation	0.54	0.07

Table 3: Top features correlated with Bad-Beat plays. Voice energy, motion energy, and voice pitch variation (std. dev. of the fundamental formant frequency) were the most correlated features.

Many, but not all of the bluffing/all-in/bad-beat situations described above were highly stressful situations. We define a stressful hand as any in which the player subjectively rated his or her stress to 7 or above (from a scale of 1 to 10). For this measure, a number of the same features described above end up being highly correlated. Skin conductance features, with the number of large-threshold peaks end up being the best feature. It was observed that many hands had small ramps of skin conductance, but large ramps of skin conductance almost always corresponded to highly stressful events. The standard deviation in the largest autocorrelation peak (a measure of pitch variation in the voice) was also highly correlated. The LF/HF ratio shows high sympathetic arousal in the heart rate variability for stressful situations.

Stress Top Features (n=59)	Statistics	
	R	p
1) Skin Conductance Peaks	0.83	0.004
2) Std. Dev. Largest Auto. Peak	0.76	0.011
3) LF/HF Ratio	0.72	0.023

Table 4: Top features correlated with stress. The number of skin conductance peaks, the std. dev. in the location of the largest auto correlation peak, and the LF/HF ratio heart rate variability measure were the most correlated features.

5. Stress and Bluffing Classification

The results in the previous section indicate that there were a variety of skin conductance, heart rate variability, voice, and movement features that can be correlated to various stressful outcomes. The question to ask is if these simple physiology features can be used to create accurate classifiers that can discriminate between these outcomes. Of particular interest in a more general setting outside of poker tournaments is stress classification, as well as potentially being able to identify when a person is lying (bluffing).

The leave-one-out cross-validated two-class linear decision classifier for stress has around 79% accuracy using skin conductance peaks. Using two features (skin conductance peaks and voice pitch variation), this accuracy goes up to around 82% ($\beta_1 = 3.5$, $\beta_2 = 1.9$). For the two-feature stress model, the Type-I errors (false positives) and Type-II errors (false negatives) represent 21% and 79% of the total misclassifications, respectively. The classification accuracies for bluffing were a less than those for stress but still significantly above chance. The cross-validated accuracy for a bluffing detector was 64% using skin conductance peaks, and 71% using both skin conductance peaks and voice energy of ($\beta_1 = 7.3$, $\beta_2 = -4.8$). For the two-feature bluffing model, the Type-I errors and Type-II errors represent 38% and 62% of the total misclassifications, respectively.

Classifier	Accuracy	
	Top Feature	2 Features
1) Stress ($\geq 7/10$)	79%	82%
2) Bluffing (Large)	64%	71%

Table 5: Classifier results for stress and bluffing. We can predict high stress events (defined as a subjective rating of 7 for stress on a 10-point scale) with 82% accuracies and bluffing (defined when the subject indicated a large scale bluff) to within 71% accuracies.

These results seem to conform to the literature in polygraphy (lie-detection techniques) [Podlesny & Raskin 1977]. The various physiological features were well understood to correlate well to stress through the autonomic nervous system reactions. Lying was capable of being detected insofar as a person becomes stressed at having to lie. In the case of poker, there was obviously a high correlation between the two (around $R=0.75$), as the player was motivated to prevent the potential loss of money. However, for the levels of play in the tournaments that were conducted, it was possible that bluffing becomes something that people adjust to since it was a relatively common occurrence.

6. Conclusions

These initial results indicate that it is possible to correlate stress, lying (in the context of bluffing), and interest with a variety of physiological features, at least within the context of structured poker settings. Although some of these correlations are no doubt idiosyncratic to the game of poker, it is reasonable to assume that these features may be applicable to more general domains, such as occupational stress. Based on the initial research, we show that it is reasonably easy to develop simple linear classifiers that can accurately predict stress based on short snapshots of physiology.

Using simple linear classifiers, we have been able to identify high stress situations to within about 82% accuracy. We can even detect lying with about 71% accuracy. Essentially, we demonstrate that we can identify these events from simple aggregated physiological features acquired during the duration of the events in question from non-invasively derived sensing. While this leaves room for improvement, the results indicate the possibility of developing classifiers that may be practical in certain domains for triaging people in stressful situations. At the very least, we can create systems that can identify individuals with a relatively high probability of being stressed out, who can then be monitored more closely for potential problems. Specifically within the domain of poker, the ability to identify stress or bluffing even slightly beyond chance provides a compelling advantage.

While the results of the study are very promising, there might have been some limitations to the methodologies we employed to find the appropriate features used in the analysis. Specifically speaking to the HRV measures, they are sometimes derived from windows under less than a minute in duration given the length of the hand (most studies involve between 3-5 minutes of recording due to errors in HRV calculations). For the voice features, many of the poker interchanges are so short in duration (shorts of people saying “call”, “raise”, etc.) that the HMM segmentation algorithm is not able to catch the voicing segments, or construct speaking segments out of them. Also, the classifiers require features that are computed on the physiology data after the hand is complete, which is not amenable to real-time classifiers which can detect stress or bluffing at the precise moment these situations occur.

The analysis in this study focuses on the outcomes correlated with stress. As we have collected a variety of other tournament information, such as interest levels and objective outcomes like money wagered, it is a potentially interesting exercise to understand how these variables correlate to physiology as well. Another future area of exploration is to use the collected physiology and behavioral data to understand the patterns of arousal and other features to see what precisely is re-enforcing addictive behavior and to better understand addiction overall. Gambling in this country is becoming more and more rampant, and in recent years poker in particular has seen a surge in popularity. If we can identify physiological markers for interest, we can potentially begin to develop a physiological model for gambling addiction.

7. References

- [Caro 2003] M. Caro (2003) *Caro's Book of Poker Tells*, Car-doza Publishing
- [Lo & Repin 2002] A. Lo & D. Repin (2002) "The Psycho-physiology of Real-Time Financial Risk Processing", *Journal of Cognitive Neuroscience*, 14:323-339
- [Lo et al. 2005] A. Lo, D. Repin, and B. Steenbarger (2005) "Fear and Greed in Financial Markets: A Clinical Study of Day-Traders", *American Economic Review*
- [Madan et al. 2004] A. Madan, R. Caneel, and Pentland (2004) "GroupMedia: Distributed Multimodal Interfaces", *International Conference on MultiModal Interfaces*
- [Marci 2002] C. Marci (2002) "Psychophysiological Correlates of Empathy", Abstract, Harvard Mysell Psychiatry Research Day, April, 2002
- [Picard 2001] R. Picard (2001) "Affective Medicine: Technol-ogy with Emotional Intelligence", Chapter in *Future of Health Technology*, IOS Press
- [Picard et al. 2001] R. Picard, E. Vyzas, and J. Healey (2001) "Toward Machine Emotional Intelligence: Analysis of Affective Physiological State", *IEEE Transactions Pattern Analysis and Machine Intelligence*, 23:10
- [Podlesny & Raskin 1977] J. Podlesny & D. Raskin (1977) "Physiological Measures and the Detection of Deception", *Psychol. Bull.*, 84(4):782-799
- [Steenbarger 2002] B. Steenbarger (2002) *The Psychology of Trading: Tools and Techniques for Minding the Markets*, John Wiley & Sons, Hoboken, NJ
- [Sung 2005] M. Sung (2005) "Non-Invasive Wearable Sensing Systems for Continuous Health Monitoring and Long-Term Behavior Modeling", Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science