

RESEARCH STATEMENT

Daniel McDuff

djmcduff@media.mit.edu

<http://alumni.media.mit.edu/~djmcduff>

There are many ways in which human-computer (and human-human) interactions could benefit from greater understanding of how people behave and express themselves. As an example, emotions play a key role in perception, memory, decision-making, health and wellbeing, and communication. Yet, our understanding of human affect is limited, in part because these processes may not be consciously accessible or easy to verbalize. Methods of accurately quantifying and modeling human emotions are essential if we are to understand the role that they play in our lives. The affective computing research field has contributed significantly towards these goals, building technology for quantifying emotional responses. However, much emotion research has been carried out in labs and/or performed with small populations that lack the diversity seen in real-life (e.g. only undergraduate students). Part of the reason for limiting experiments to the lab is that dedicated hardware has been required to capture and quantify responses accurately. There are inherent ecological validity questions if everything is performed in a controlled and/or sterile setting. In order to advance the fundamental understanding of human emotions, build smarter affective technology, and ultimately help people, we need to perform research in-situ and with a representative population. To do this there are challenges that need to be overcome. We need: i) Data collection frameworks that use ubiquitous sensors to capture large-scale and longitudinal observations; ii) Measurement algorithms that can quantify physiology and behaviors robustly in-situ; iii) Deployment and evaluation of applications in the real-world. **My Ph.D. thesis and subsequent work has addressed aspects of these three challenges associated with behavioral research, allowing me to collect and analyze observational data from millions of individuals across the globe.** The connectivity of the Internet provides an efficient channel for collecting data from huge populations across large geographic areas. We can also make use of the vast array of sensors available on ubiquitous devices (laptops, cellphones, wearables), the hardware is available in many areas where expensive custom hardware is not.

Past Work

Data Collection: Ecologically Valid Data Gathering at Massive Scales

In recent years crowdsourcing has proliferated widely within many areas of computer science and human-computer interaction (HCI). Affective Crowdsourcing is the application of crowdsourcing techniques to affective computing. Crowdsourcing approaches to the tasks of data collection and labeling have the potential to vastly reduce the time and cost associated with research and, more importantly, allow us to increase the scale and reach of the research considerably. In 2011 we launched the first large-scale collection of facial responses over the Internet. In one month I was able to collect over three thousand responses from thousands of unique individuals. Subsequently, I have run many more studies refining the process and collecting still larger amounts of affective data. This work was the starting point for what is currently the largest dataset of observational emotion data in the world.

Measurement: Novel Approaches for Quantifying Behavior and Physiology

1. *Automatic Detection of Facial Expressions*

The face is one of the richest sources of affective information. The facial action coding system (FACS) is the most comprehensive and widely used taxonomy for quantifying facial behavior. During my Ph.D. I used automated techniques of quantifying the facial expressions (based on FACS) within the webcam videos described above. A majority of these videos were collected in uncontrolled settings over the Internet and therefore presented challenges for computer vision analysis. In addition, I developed novel approaches for detecting facial actions from video. In this work we utilized the sparsity and co-occurrence of facial actions to gain state-of-the-art performance on benchmark spontaneous facial expression datasets.

2. *Physiological Measurement using Digital Cameras*

I developed the first method for measuring the blood volume pulse (BVP) via photoplethysmography (PPG) using just webcam. This work received a \$60K grant from the Center for Integrated Medicine Innovative Technology (CIMIT) & a Top 10 Invention of the Year Award (Popular Science). The method allows the accurate measurement of heart rate, respiration rate and heart rate variability (HRV). Remote

measurement of physiology using just a webcam has several benefits: 1) measurements can be taken in more naturalistic settings, 2) measurements can be scaled easily, 3) the hardware is very low-cost.

In collaboration with Olympus I have extended this research using a novel five-band camera (no more costly than a normal CMOS sensor) showing increased accuracy in physiological measurement using the additional color bands. Furthermore, we can detect very subtle aspects of the pulse wave morphology, capturing timing of both the systolic and diastolic peaks/inflections. Affective computing and healthcare applications can both benefit significantly from advancements in sensing capability such as this. BVP morphology, in particular, has been linked with aging and arterial stiffness. I am now working, in collaboration with the Air Force Research Lab, with multi-camera and multi-spectral arrays to further improve these methods.

3. *Physiological Measurement using a Smartwatch or Smartphone*

Visual methods of capturing physiological measurements have limitations. Once the subject is out of the field of view of the camera, or there is insufficient illumination, measurement is not possible. In recent work, we have shown that subtle body motions captured via accelerometers and gyroscopes in a wearable device (or even just a smartphone) can also allow us to measure heart rate and respiration rate. Furthermore, we can use these for identity and posture recognition, allowing for continuous and passive identity verification. This work has been awarded several best paper awards and nominations in the past year.

Deployment: Applications and Evaluation

1. *Consumer Intentions and Marketing Effectiveness*

In my Ph.D. thesis work I modeled how emotional responses to media content are related to preferences and marketing effectiveness. In the largest international ad study to date I captured over 20,000 responses to online video ads and found that ad liking and an individual's purchase intent was strongly related to their facial response. Furthermore, by collaborating with Mars, Inc. I was able to show that these measurements were correlated with sales effectiveness (the impact of the ad on short-term sales).

2. *Voters' Preferences*

During the 2012 US presidential election campaign we performed the first analysis of facial responses to presidential election debates using viewers' webcams. Analyzing the facial valence of hundreds of viewers I was able to train an accurate model for predicting candidate preferences. In addition, I was able to identify salient sequences of expressions involving subtle asymmetric smirks as well as smiles and brow furrows.

3. *Capturing Arousal*

Using our novel approach for remote measurement of physiology I demonstrated that the technique was able to capture subtle changes in autonomic nervous system (arousal) activity when subjects were under cognitive load. By analyzing videos of people completing computer tasks we were able to build a predictive model for cognitive load. This work complements the analysis described above and shows that we can capture arousal (via HRV) in addition to valence (via facial expressions) from videos of the human face.

4. *Life-logging and Emotion Visualization*

The visualization of affective data is an important part of interpretation and understanding. To explore this I developed AffectAura an emotion prosthetic that involved multi-modal sensors, an emotion prediction algorithm and an interface for user reflection. The system incorporated contextual data alongside predictions and allowed us to investigate the utility of emotion tracking in a real-life setting. The Medical Mirror and Inside-Out are other interfaces I have designed that allow users to view affective data.

Current Work

In my current research I am extending the results from my thesis work using the largest collection of emotion data in the world. I am now able to deploy the models I have trained across responses to thousands of pieces of content (~4,000,000 unique face videos). We are able to characterize how affective responses differ across cultures, age and gender. This work is the first to present findings of large-scale differences in observational measurements of facial behavior across different demographics. It provides powerful evidence that the field of affective computing has advanced to the point of enabling research that would have been impossible using traditional methods of data collection and analysis. I believe it will transform how observational analysis of human behavior is performed.

I have been developing facial action unit (AU) classifiers that work robustly in challenging environments and detect very subtle behaviors. Using the largest dataset of spontaneous and naturalistic facial responses in the world I have been able to systematically analyze the impact of different types of training data, sampling methods, features and models. Increasing the size and diversity of examples with the training set significantly improves the overall discriminability of the system for all action units considered. In this work we have developed and evaluated an efficient support vector machine (SVM) non-linear kernel (radial basis function (RBF)) approximation that shows greater accuracy for AU classification but similar computational cost to a model with a linear kernel. This has allowed us to build a set of real-time classifiers that performs extremely well on challenging real-world data. The videos that are available also allow for exploration of techniques that can leverage large amounts of unlabeled data.

Through my work at the MIT Media Lab I have collaborated with many big corporate research departments: Olympus (collaboration on camera prototype), Mars (large-scale thesis study into the role of emotions in advertising - I raised \$180,000 in funding for this project), Procter and Gamble (fellowship - 2nd year Ph.D.), NEC (fellowship - 4th year Ph.D.), Microsoft Research. These companies represent potential future funding sources for my research. I currently also have funding from the Air Force Research Lab. I have interacted and worked with many other academics within, and outside of, MIT. I have built up collaborations that will be very beneficial in the future, as my work lies at the nexus of a number of research fields (affective computing, computer vision, crowdsourcing, biomedical engineering).

Future Work

There are many examples in which human-computer (and human-human) interactions could benefit from greater understanding of affective signals. In my future work I will focus on how automated facial and physiological measurement can be integrated into real-world applications on a large-scale. Specifically, I am interested in how we can take scaled emotion measurement beyond understanding an individual's response to video and advertising content, which was the focus of my thesis work. I aim to research how emotions can be used to improve communication and wellbeing. How can emotions be measured during online learning for improving communication and teaching? How can we measure wellbeing by quantifying affect during everyday tasks (e.g. remotely capturing physiology during computer-based work)? How is affective data best represented and communicated through digital interfaces? I will design experiments in which we can obtain longitudinal observations (capturing affective responses from the same individuals across time) through inputs from multiple modalities (video, motion, audio) available via ubiquitous sensors. These research goals build on my current work and open up a wealth of new research directions.

Achieving this end will require further development of remote physiological sensing and facial expression measurement. I aim to continue to improve automated facial expression analysis with the help of the huge amounts of data I collected during my Ph.D.. I will continue to develop physiological measurement techniques using low-cost cameras. There are algorithmic improvements that will allow these methods to work robustly over bandwidth-limited connections and with the subject in a variety of lighting conditions. Furthermore, my research into multi-spectral imaging techniques for physiological measurement has enormous potential. I believe these tools could also benefit other researchers within the university and setup cross-disciplinary collaborations.

However, my work will not be limited to improving the accuracy of remote facial expression and physiological sensing. The deployment of technology in the real-world is vital in order to assess its impact and how it will benefit people. For affective data to be useful in real-world applications, interfaces and visualizations need to represent and communicate the insights in a suitable way. I will continue my research into digital interfaces for communicating affective information. These interfaces need to be tested through rigorous user-studies.

Finally, an area in which emotion measurements can play an important role is in our understanding of human wellbeing. A huge potential for my research is in measuring affective cues at scale in order to understand the links between our emotions, lifestyle and health. The advances in sensing technology that I have presented can help make this possible. I plan to develop cross-disciplinary research projects that address problems associated with health and wellbeing. My work intersects the fields of affective computing, computer vision, crowdsourcing and biomedical engineering providing great potential for collaboration across the department, university and with sponsors. During my Ph.D., I found that collaborations with faculty and students were mutually beneficial and extremely rewarding.