

Modeling Humans via Physiological and Behavioral Signals

Insights

- Physiological and behavioral signals such as galvanic skin response or simple finger gestures on a mobile phone can help model humans.
- A generic framework linking stimuli to models using machine learning promises objective and reliable measurements.
- Two practical case studies illustrate the framework's potential for detecting psychological traits and early diagnosis of Parkinson's disease.

The time has come to accurately and unobtrusively model humans! Modeling humans—whether in terms of their skills, emotions, or attitudes—can help us deliver tailored services, interaction, and information. Let's teach math by associating each formula with other concepts and formulas the student already knows. Let's recommend movies based on the emotions that past movies elicited in an individual. Let's produce data visualizations in line with what has previously resonated with a user. Why are these objectives so obviously needed, yet so elusive? A key reason is that a lack of integration between

tasks, human behavior, and reasoning is hindering the full power of interaction customization. Modeling humans can help bring these elements together.

In recent decades, HCI and AI have developed a range of tools and methodologies for modeling humans and personalizing interaction [1], but these often remain bounded by the information needed by the application itself. For example, movie recommenders generally focus on behavioral signals, such as past movie selections by the user or other users of the same platform, sometimes looking at related factors like social



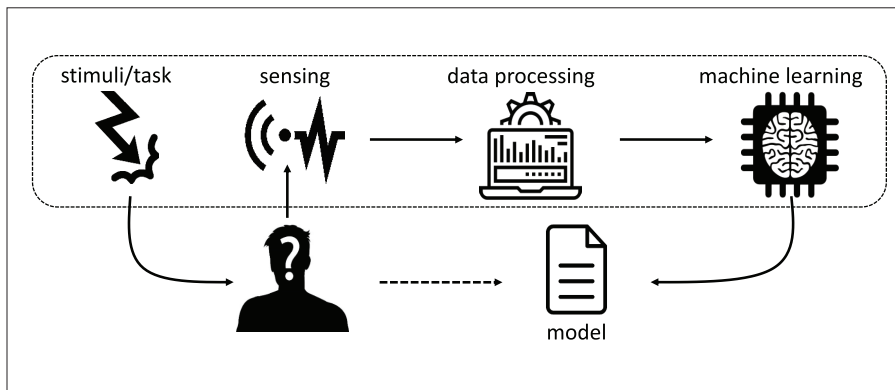


Figure 1. Human-modeling framework.

media likes or movie reviews. Human modeling, however, rarely ventures into gauging thoughts, emotions, or attitudes, at best resorting to clunky pop-ups such as “Would you prefer option A or B?”

Another common drawback of many human modeling tools is that they can be manipulated, making their reliability questionable. For example, it is easy to pose as a horror-movie lover in a system, or similarly to spread such fake information on social media. In fact, unbalanced training data will result in strong biases even without any manipulations [2]. Models that capitalize on direct input are even easier to trick, as the human user may have a good idea about the desired answers, allowing them to steer the system in the right direction. These risks intensify in high-stakes scenarios, such as performance evaluation or job recruitment. Here, humans may be willing to paint a fictitious picture of certain behavior, overstate their skills and knowledge, or simply provide inaccurate information, which may hinder the modeling and affect its outcomes.

A solution is at hand; it is all a question of piecing things together. In order to make human modeling more objective and reliable, let’s turn to the shelves of physiologists, neuroscientists, and sensing engineers. A multitude of

physiological and behavioral signals generated by the human (consider heartbeat, brain activity, skin conductance, blood pressure, and eye movements) can hardly be consciously controlled and can potentially disclose precious and reliable information about the human experience, if properly harnessed. While this raises significant technical challenges, requiring a combination of sensing, signal processing, and machine-learning skillsets, it also has a tremendous potential to pave the way for next-generation human-modeling methods.

A HUMAN-MODELING FRAMEWORK

Based on our recent work on the detection of personality traits [3], as well as another work on the prediction of Parkinson’s disease [4] (both published at CHI 2019), we present here a framework for such physiological and behavioral signal-based human modeling (Figure 1).

The main idea behind the framework is that consciously uncontrollable signals can be treated as objective predictors for the model. To offer a reliable and measurable input, such signals need to come in response to a standardized *stimulus* or *task* triggering them. A range of triggers can be used—from passive exposure to multimedia content to

highly interactive activities—as long as the triggers are causally linked to the derived model. Domain expertise is required for this step. For example, when teaching mathematics, such a trigger could be a series of mathematical exercises of increasing difficulty. The modalities of the stimuli and tasks are crucial, as they are likely to trigger different physiological signals. For example, as our work discussed below shows, signals elicited by a video clip would be stronger than those elicited by an image.

The physiological signals are captured by a *sensing technology*. Guided by the main goal of objective human modeling, we prioritize technologies that can capture difficult-to-control signals. Indeed, some physiological or behavioral signals, like breathing rate or eye gaze, can be consciously controlled; others, like blinking, blood pressure, or heart rate, are difficult to control; while many other signals, like skin conductance or electric brain activity, cannot be controlled at all. The level of control that humans exhibit over the signal drives the objectivity of the captured data and the reliability of the derived models. Another consideration is the practicality and ease of use of the sensors, which ideally should not impede normal behavior and interactions, and should mitigate artifacts related to movement, temperature, lighting, and more.

While it may be appealing to bring in many sensors, they help only once useful features are extracted; otherwise they are simply flooding the systems with irrelevant data. Raw sensor data, like skin-conductance values, electric brain signals, or heart-rate data needs to be preprocessed, and its statistical characteristics extracted, during *data processing*. While this varies according to the deployed sensor, three typical preprocessing steps are: filtering (data cleansing, noise reduction, and artifact removal), segmentation (partitioning into time intervals), and normalization (with respect to a baseline signal). These are followed by feature extraction, which again depends on the sensor and the stimuli. For example, analyzing the shape of spikes in electrodermal activity may

The level of control that humans exhibit over the signal drives the objectivity of the captured data and the reliability of the derived models.

reflect cognitive responses to a short math question but may be less useful when aggregated over the duration of a feature-length movie.

Finally, the processed data can be used to build the human models using *machine learning*. The model consists in a set of parameters trained from labeled data—features extracted from physiological signals of humans with already established models—which serve as ground truth examples (marked by the dashed arrow in Figure 1). These labels are collected using traditional methods such as questionnaires or observations, and hence are prone to noise and manipulation. However, system developers can control quality through incentives (payment for truthful responses) or disincentives (no benefit from cheating). A range of machine-learning algorithms are readily available to train the human-modeling component of the framework. Once trained, it can be deployed to predict the model values for new subjects, whose features are also extracted from their physiological signals. When too many features are extracted, machine learning is at risk of overfitting, but this can be mitigated using feature-selection methods.

TWO RECENT HUMAN-MODELING CASE STUDIES

W*We predicted human personality traits using affective stimuli and eye-tracking data [3].* Personality is considered to be a set of stable characteristics that affect human behavior, cognition, and emotions. Personality detection is a nontrivial task, typically requiring humans to fill out lengthy questionnaires rooted in personality and psychometric theories. Since the questionnaire data is self-reported, the models are often noisy and manipulation-prone.

We showed our participants 50 images and seven videos, all validated to evoke emotional responses. We focused on eye signals captured using commercial-grade eye-tracking glasses—eye blinks, saccades, fixations, and pupil-size measurements. Ten features were extracted, ranging from simple ones like saccade rate per second, to more complex geometric ones intended to

reflect ocular muscle activity, such as average of the peak angular velocity of each saccade. We extracted a total of 170 features from the images and videos. Since the data collection involved only 21 participants, we applied correlation-based feature selection to select a predictive set of fewer than 10 features. These were fed into machine-learning classifiers trained to predict 16 personality traits across three established personality models: Dark Triad, BIS/BAS, and HEXACO [5]. The ground truth data was obtained by administering the personality questionnaires associated with these models, and grouping the participants into low, medium, and high classes for each trait.

The overall predictive accuracy of the best-performing classifier, across all the personality traits and participants, was close to 0.86, with six traits being predicted with an accuracy greater than 0.90. This is much higher than the benchmark random-guess probability of 0.33 in the three-class classification task. In particular, we noted that the tactics, views, and morality traits achieved over 0.90 accuracy; all belong to the Machiavellianism component of the Dark Triad, which is associated with affective rather than cognitive traits [6]. We attributed this to the fact that our stimuli were affect-based. Considering the image and video stimuli, we unsurprisingly found that the videos resulted in more accurate predictions than the images. Similarly, we examined the predictive power of features for various traits and found the most predictive feature-trait combination was the number of blinks and psychopathy, which aligns with prior work showing that people with psychopathic traits tended to display unusual blink responses [7].

Practically speaking, such a system could dramatically reduce the time required to administer questionnaires and provide near real-time objective personality modeling. Accuracy rates of 0.86 may not yet allow the detection of mental pathologies but could prove useful in longitudinal psychological assessments. The low entry cost of sensors and data-analysis packages suggests that such a method could practically

supplement traditional personality questionnaires.

Predictions of Parkinson's disease (PD) from mobile phone gesture analysis were achieved using a similar approach [4]. PD is a neurodegenerative disorder that affects the motor system and is diagnosed using neurological examinations, computer tomography scans, and magnetic resonance imaging. It is important to highlight that PD symptoms often show up—and, thus, PD is diagnosed—at relatively late stages, when significant and irreversible brain damage has already occurred, emphasizing the importance of early detection.

To predict the PD diagnosis, the authors used commercial-grade smartphones and common mobile gestures: flick, drag, handwriting, pinch, and tapping. Since PD affects humans' fine-motor skills [8], the authors hypothesized that this will manifest in finger movements captured by the smartphone's sensors. The participants were tasked with performing 60 flick, 60 pinch, and 30 drag gestures, writing and typing for 10 minutes, and performing the alternative finger-tapping test currently used to diagnose PD. The touch signal captured by the screen sensors was processed and 46 features extracted, which were grouped into touch, trajectory, temporal, and inertia groups. The study involved 102 participants: about a third diagnosed with PD and the rest healthy. Hence, the human modeling was essentially a PD diagnosis prediction represented by a two-class classification.

The accuracy of the predictions was measured using the area under the receiver-operating characteristic curve (AUC) that ranges from 0 to 1. First, the authors studied the performance of the four groups of features. Although trajectory and inertia features achieved AUC close to 0.9, combining all the groups improved the AUC to 0.95. Considering the mobile-input gestures, drags and pinches achieved an AUC of 0.92. Further adding flicks and handwriting boosted the AUC to 0.95, while eventually adding typing increased the AUC to as high as 0.97. Overall, for the best-performing combination of gestures and features, both true positive and true negative

rates—that is, the ratio of correct predictions for PD-diagnosed and healthy subjects—were around 0.9. To position this with respect to existing clinical methods, the currently deployed alternative finger-tapping test achieves AUC of approximately 0.83 [9].

Practically speaking, this case study showcased another implementation of the proposed framework in a promising medical application of human modeling, using a simple smartphone. By joining cross-disciplinary skillsets, the authors demonstrated that a challenging medical condition can be predicted with accuracy levels surpassing the current clinical methods. Research is yet to study whether neurological conditions can be detected with electro- and magneto-encephalogram (EEG and MEG, respectively) sensors, directly capturing brain activity.

WHERE TO NEXT?

With recent advances on the sensing, signal/data processing, and machine-learning fronts, the exercise of accurate and reliable human modeling seems to be within reach. The modeling framework we proposed seeks to provide structure and confidence to teams aiming to boost user experience by introducing human modeling and personalization in their applications. The discussed case studies are promising; in the near future, we may be able to:

- **Predict mental/cognitive disorders.** With a plethora of new body and activity-tracking devices, mental or cognitive disorders could be assessed on a high-frequency basis by the proposed framework, instead of requiring a visit to a practitioner's clinic. Conditions such as anxiety and depression in the mental space, or even reading and learning disorders like dyslexia, are often hard to establish objectively but could potentially be screened using our framework and the right combination of stimuli, sensors, data processing, and machine learning. Affective stimuli were shown to be accurate in the first case study discussed here; hence, suitable cognitive triggers could

also be developed to target specific disorders, which can potentially be captured by common sensors such as a camera/microphone or motion trackers. Stimulus-based processing also safeguards against misuse of the system, as the evaluation takes place in an agreed time and place. Such screening technologies would be invaluable for effective cognitive-behavioral therapies.

- **Detect susceptibility to cybersecurity attacks.** The human factor plays a major role in cybersecurity, and with the existing security software in place, many cyber incidents are now associated with human error. For example, making a hasty decision on an incoming email, potentially a phishing attack, can have disastrous consequences for the targeted individual and their organization alike. Hence, it is critical to understand who is susceptible to what type of cybersecurity risks, to be able to educate and protect these people. Deploying our framework in order to examine the brain signals or mouse movements of a person faced with simulated cybersecurity threats could allow the modeling of their cognitive processes. Once such a model is derived, it will be possible to detect human hesitation or subconscious behavior when faced with potential cyber threats, and bring this to their attention for conscious examination. Combining the model with adaptive training could further reduce vulnerability and upgrade human users into an active defence against cyberattacks.

We believe the HCI community and available technology now provide the required support for novel next-generation methods for human modeling and personalized interactions. The use cases we presented highlight how important real-life problems have been addressed with promising results, and can be abstracted in a reasonably simple framework. There are many more challenging problems out there offering high-reward and real-life impact. Hence, we are calling for action and take this opportunity to encourage researchers and practitioners to look into these problems!

ENDNOTES

1. Kobsa, A., Nejdl, W., and Brusilovsky, P., eds. *The Adaptive Web: Methods and Strategies of Web Personalization*. Springer, 2007.
2. Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proc. of the National Academy of Sciences* 116, 32 (2019), 15849–15854.
3. Berkovsky, S., Taib, R., Koprinska, I., Wang, E., Zeng, Y., Li, J., and Kleitman, S. Detecting personality traits using eye-tracking data. *Proc. of CHI 2019*, paper 221.
4. Tian, F., Fan, X., Fan, J., Zhu, Y., Gao, J., Wang, D., Bi, X., and Wang, H. What can gestures tell? Detecting motor impairment in early Parkinson's from common touch gestural interaction. *Proc. of CHI 2019*, paper 83.
5. McCrae, R.R. and Costa, P.T. *Personality in Adulthood: A Five-Factor Theory Perspective*. Guilford Press, 2003.
6. Jonason, P.K. and Krause, L. The emotional deficits associated with the Dark Triad traits: Cognitive empathy, affective empathy and alexithymia. *Personality and Individual Differences* 55, 5 (2013), 532–537.
7. Patrick, C.J., Bradley, M.M., and Lang, P.J. Emotion in the criminal psychopath: Startle reflex modulation. *Journal of Abnormal Psychology* 102, 1 (1993), 82–92.
8. Pradhan, S., Brewer, B., Carvell, G., Sparto, P., Delitto, A., and Matsuoka, Y. Assessment of fine motor control in individuals with Parkinson's disease using force tracking with a secondary cognitive task. *Journal of Neurologic Physical Therapy* 34, 1 (2010), 32–40.
9. Arroyo-Gallego, T., Ledesma-Carbayo, M.J., Butterworth, I., Matarazzo, M., Montero-Escribano, P., Puertas-Martin, V., Gray, M.L., Giancardo, L., and Sánchez-Ferro, A. Detecting motor impairment in early Parkinson's disease via natural typing interaction with keyboards: Validation of the neuroQWERTY approach in an uncontrolled at-home setting. *Journal of Medical Internet Research* 20, 3 (2018), e89.

👤 **Ronnie Taib** is a principal research engineer at Data61 – CSIRO in Sydney, Australia. With a passion for understanding and measuring human-machine interaction, he has published over 50 papers covering multimodal interaction and cognitive load measurement based on physiology and behavioral signals.
→ ronnie.taib@data61.csiro.au

👤 **Shlomo Berkovsky** is an associate professor at the Australian Institute of Health Innovation, Macquarie University, where he is leading a team of researchers working on precision health. He is a computer scientist who has published over 130 papers. His core expertise areas are user modeling and personalized technologies.
→ shlomo.berkovsky@mq.edu.au