

Cyber Physical System Challenges for Human-in-the-Loop Control

Sirajum Munir, John A. Stankovic
University of Virginia
{munir, stankovic}@cs.virginia.edu

Chieh-Jan Mike Liang
Microsoft Research Asia
liang.mike@microsoft.com

Shan Lin
Temple University
slin@temple.edu

Abstract

This paper articulates three main challenges for employing feedback control with humans in the loop. They are: (i) the need for a comprehensive understanding of the complete spectrum of the types of human-in-the-loop controls, (ii) the need for extensions to system identification or other techniques to derive models of human behaviors, and (iii) most importantly, determining how to incorporate human behavior models into the formal methodology of feedback control.

1 Introduction

Human-in-the-loop feedback control systems offer exciting opportunities to a broad range of cyber-physical system applications including energy management [7], health care [5], and automobile systems [2]. For example, it is hypothesized that explicitly incorporating human-in-the-loop models for driving can improve safety, and using models of activities of daily living in home health care can improve medical conditions of the elderly. Although having humans in the loop has its advantage, modeling human behaviors is extremely challenging due to the complex physiological, psychological and behavioral aspect of human beings. Here we propose that it is necessary to raise human-in-the-loop control to a central principle in system design and to solve three main challenges.

2 Challenges

Challenge 1: *The need for a comprehensive understanding of the complete spectrum of types of human-in-the-loop controls.*

There are many variations for human-in-the-loop controls. We need to understand the complete spectrum to determine the underlying principles and subtleties that separate them. Figuring out the principles will be very

beneficial for the future applications and allow the creation of feedback control solution techniques for each of them, building on commonalities when appropriate.

To begin to understand the spectrum of human-in-the-loop controls, we start with creating a taxonomy of human-in-the-loop applications based on the controls that they employ. We classify the applications into three categories: (i) applications where humans directly control the system, (ii) applications where the system passively monitors humans and takes appropriate actions, and (iii) a hybrid of (i) and (ii). Now, we describe some example applications belonging to each of these categories.

Applications where humans directly control the system primarily use supervisory control. In supervisory control, involvement of humans takes place in two ways. In one case, the process runs autonomously. Humans intervene with the control algorithm when it is necessary typically by adjusting set points. These control problems are well understood. In the other case, the process accepts a command, carries out the command autonomously, reports the results and waits for further commands to be received from the human. As a concrete example, in [10], human-in-the-loop control is used in a wheelchair-mounted robotic arm to retrieve an object from a shelf. In this feedback control system, human provides input via a touch screen or joystick which is analyzed by a vision processing system to position the robotic arm to retrieve the object. In this application, a human directly controls the controller of the feedback control system and guides it to take appropriate action. The main problems here involve interfacing humans to control loops.

A broad range of applications monitor human behavior passively and take appropriate actions (including doing nothing). These applications are sub-categorized to either open loop or close loop system. For example, Lullaby [5] uses a commercial off the shelf sleep tracking device to keep track of sleep quality. It also uses sound,

light, temperature, and motion sensors to record the environmental conditions during sleep. All this information is presented to the users in a tablet that helps them to identify the potential causes of sleep disruption. Here, human is in the loop, but the human does not directly control the system and this is an open loop system as Lullaby does not proactively take any action to improve the quality of sleep. Similarly, another open loop system is AlarmNet [11], which is a state of the art assisted-living and residential monitoring network for smart home health care. It monitors activities of daily living by using the environmental and wearable sensors and creates a continuous medical history. Authorized health care providers are allowed to monitor activity patterns to determine if the residents need immediate attention or new healthcare services. Open problems here include understanding what human behaviors need to be monitored.

On the other hand, Smart Thermostat [7] is a closed loop human-in-the-loop system as it uses sensors to detect occupancy and sleep patterns in a home and uses these patterns to proactively turn off the HVAC system to save energy. Another example of closed loop system is an advanced driver assistance system. In [6], a driver's intended actions, e.g., to turn or lane change are inferred from several sources, including the driver's current control actions (steering and acceleration), his visual scanning behavior, and the surrounding traffic environment. Although [6] is an open loop system, if we can model the driver's physiological state (fatigue, anger, drunk, etc.) and behaviors (distraction, erratic steering, etc.), then we can convert it to a closed loop system where the automobile can immediately react and signal alarms, or even wrestle control from the driver when the driver is unfit to keep the safety or fuel efficiency of the current trip. In these systems, real-time response and intimately incorporating physiological, behavioral and psychological aspects of humans are complex open questions.

The hybrid system passively monitors human behavior, takes appropriate actions and also takes occasional human inputs for the control. For example, in [12], human-in-the-loop is used to control building energy. It uses human feedback to adjust set point of the control, e.g., human feedback is used to adjust the temperature set point of the HVAC system to maintain thermal comfort. It tracks the position of the occupant and feeds this information to the controller so that energy is delivered to only those spots where needed. These systems incorporate all the challenges from the above discussion. After understanding these behavioral aspects, our next challenge addresses how to model these behaviors using the appropriate modeling techniques.

Challenge 2: *The need for extensions to system identification or other techniques to derive models of human behaviors.*

System identification is a powerful technique to create system models. It is a new challenge to apply it to human behaviors. The order and types of equations to use, how to produce adequate testing inputs, what output variables are required, and how such a model accounts for human traits are unknown. For example, Empath [3] is a real-time depression monitoring system for the home. It collects different behavioral data including sleep, weight, activities of daily living, and speech prosody and can potentially detect the early signs of a depression episode, as well as can track progress in managing a depressive illness. This is an open loop system as the system only shows the reports to the caregivers and does not proactively take any action to improve the quality of life of the depressed patients. If we were to use system identification technique to model a human being who is suffering from depressive illness, it is not clear what are the inputs, what are the states and how the state transitions occur based on different physiological, psychological and environmental factors. If there was a formal model of human behavior or even an estimated model, then by combining all the factors that affect depression, we could close the loop by changing the factors in a way that helps the patients and that is based on an established methodology rather than ad hoc rules.

Capturing human behavior by extending system identification or other modeling techniques is extremely difficult due to complex physiological, psychological and behavioral aspects of human beings. Also, the level of modeling depends on application requirements. Although requirements are different for different applications, a significant portion of human-in-the-loop applications have to address some common challenges, e.g., user specific thresholds and parameters, change of human behavior over time, and required sensing technology to sense the appropriate aspects of human behavior. We need to model human behavior for large number of applications before general principles and theories emerge to address these issues. Clustering, data mining, inference, first principle models based on human physiology and behaviors may all be necessary techniques to be enhanced and applied for different applications. Robust CPS systems will likely require predictive models to avoid problems before they occur, consequently advances to stochastic model predictive control are also required. It is also unlikely that any models developed initially to design the controllers will remain accurate as the system and human behaviors evolve over time. Hence, adaptive control with humans-in-the-loop will be necessary.

Currently, state of the art techniques that model certain aspect of human behavior are either very general or very specific. For example, Smart Thermostat [7] uses a Hidden Markov Model (HMM) to model occupancy and

sleep patterns of the residents in a home to save energy, which captures human behavior from a very high level. On the other hand, [4] proposes mathematical models for impulsive injection of insulin for diabetes mellitus. Their model determines the insulin injection by closely monitoring the glucose level when it reaches a threshold value. It addresses some critical challenges for designing an artificial pancreas. In another case, [8] models human muscle in a closed loop system based on high order sliding mode techniques. The controller is used to control shank movement and has shown a great accuracy and robustness against force perturbation. It can be useful in Functional Electrical Stimulation (FES) for the paraplegic patients in regaining limited locomotor activities through electrical stimulation of the lower extremity muscles. As another example, [9] describes a new paradigm called Body Coupled Communication (BCC) for Wireless Body Area Network (WBAN) that leverages the human body as a communication channel. In this application, sensors are implanted in a human body that are capable of monitoring a wide range of physiological and emotional states and human serves as a communication channel to transmit the sampled data to a centralized monitoring entity. Here, human is modeled as a body channel where different layers of the body, e.g., skin, fat, muscles and bone offer varying but measurable levels of impedance to the signal. The human modeling in the latter three applications are very specific to the application.

While understanding of human physiological responses will continue to improve and expand to new problem areas, we also need to understand holistic human behaviors. To achieve these goals, we need to advance the state of the art in system modeling techniques.

Challenge 3: *Determining how to incorporate human behavior models into the formal methodology of feedback control.*

Even if we have a model of human behavior, it is not clear where to place the model for “each” application. For example, consider the Smart Thermostat [7] project where the system monitors human occupancy and turns off the HVAC when not needed. In this case, humans are in the loop, but they are not giving any active feedback. Now assume that sometimes humans change the temperature set point. They are still in the loop, but they are part of a hierarchical control where they are controlling some parameters from the higher level while in the lower level the system is trying to reach the set point. Now assume a more sophisticated system where the room temperature is being changed depending on physiological or psychological status of the human, e.g., when someone is experiencing an episode of depression, the room temperature is increased. In such a system, the human state as detected by sensors acts to guide the control system when a depression episode begins. Apart from this context, in

general, there are several areas where a human model can be placed:

- Outside the loop,
- Inside the controller,
- Inside the system model,
- Inside a transducer, and
- At various levels in hierarchical control.

The newest challenge seems to be how to incorporate the human behavior **as part of the system itself**. Can we define/guarantee/learn the stability, accuracy, settling time and overshoot properties of such systems, initially and as the system and human behavior evolves? As an example, [1] proposes a procedure to refine user behavior models based on reports of accidents and incidents that occur during the operation of electrical power system. This work mainly focuses on using Components Model of Emotion (CME) for observing, recording and analyzing the emotional components of the operator behavior, which can be eventually useful for simulating dynamic behavior of an operator performing tasks in a context that leads to an error. If we can model such an operator behavior using formal methodology of feedback control and if we can incorporate these operator models into the system, we will be able to analyze various safety properties of the overall system.

Incorporating human models into the formal methodology of feedback control has several advantages. For example, it allows us to analyze the property of the whole system using existing or new feedback control strategies. Also, we can run optimization techniques across people in a home, or in a building, or in a city considering various metrics and choose optimal parameters to maximize multidimensional utilities, e.g., health improvement as well as saving energy.

3 Summary

In this challenge paper, we identified three major research challenges of cyber physical systems involving human-in-the-loop control. The challenges are: (i) understanding the complete spectrum of human-in-the-loop control since more sophisticated human-in-the-loop applications are appearing, (ii) modeling human behavior of various types and identifying the best modeling schemes for each type, and (iii) incorporating these models into the formal feedback control methodology which may require new results and theory to support formal performance guarantees. Basically, state of the art system modeling techniques and feedback control strategies need to be advanced to address these challenges.

4 Acknowledgement

This work has been funded, in part, by NSF grants CNS-1035771 and CNS-1239483.

References

- [1] AGUIAR, Y., VIEIRA, M., GALY, E., MERCANTINI, J. M., AND SANTONI, C. Refining a user behaviour model based on the observation of emotional states. In *COGNITIVE* (2011).
- [2] BURNHAM, G. O., SEO, J., AND BEKEY, G. A. Identification of human driver models in car following. *IEEE Transactions on Automatic Control* 19, 6 (1974), 911–915.
- [3] DICKERSON, R. F., GORLIN, E. I., AND STANKOVIC, J. A. Empath: a continuous remote emotional health monitoring system for depressive illness. In *WH* (2011).
- [4] HUANG, M., LI, J., SONG, X., AND GUO, H. Modeling impulsive injections of insulin: Towards artificial pancreas. *SIAM Journal of Applied Mathematics* 72, 5 (2012), 1524–1548.
- [5] KAY, M., CHOE, E. K., SHEPHERD, J., GREENSTEIN, B., WATSON, N., CONSOLVO, S., AND KIENZT, J. A. Lullaby: a capture & access system for understanding the sleep environment. In *UbiComp* (2012).
- [6] LIU, A., AND SALVUCCI, D. Modeling and prediction of human driver behavior. In *Intl. Conference on HCI* (2001).
- [7] LU, J., SOOKOOR, T., SRINIVASAN, V., GAO, G., HOLBEN, B., STANKOVIC, J., FIELD, E., AND WHITEHOUSE, K. The smart thermostat: using occupancy sensors to save energy in homes. In *SenSys* (2010).
- [8] MOHAMMED, S., FRAISSE, P., GUIRAUD, D., POIGNET, P., AND MAKSSOUD, H. Towards a co-contraction muscle control strategy for paraplegics. In *CDC-ECC* (2005).
- [9] SCHIRNER, G., ERDOGMUS, D., CHOWDHURY, K., AND PADIR, T. The future of human-in-the-loop cyber-physical systems. *Computer* 46, 1 (2013), 36–45.
- [10] TSUI, K. M., KIM, D.-J., BEHAL, A., KONTAK, D., AND YANCO, H. A. I Want That : Human-in-the-Loop Control of a Wheelchair-Mounted Robotic Arm. *Journal of Applied Bionics and Biomechanics* 8 (2011).
- [11] WOOD, A., STANKOVIC, J., VIRONE, G., SELAVO, L., HE, Z., CAO, Q., DOAN, T., WU, Y., FANG, L., AND STOLERU, R. Context-Aware Wireless Sensor Networks for Assisted Living and Residential Monitoring. *IEEE Network* 22, 4 (Jul./Aug. 2008), 26–33.
- [12] ZEILER, W., HOUTEN, R. V., BOXEM, G., VISSERS, D., AND MAAIJEN, R. Indoor air quality and thermal comfort strategies: the human-in-the-loop approach. In *ICEBO* (2011).