When, and to what, does human-machine coadaptation converge?

Benjamin Chasnov, Lillian J. Ratliff, Samuel A. Burden
{bchasnov,ratliffl,sburden}@uw.edu
Electrical and Computer Engineering, University of Washington, Seattle, WA, USA

I. Background

Machines are designed to optimize a performance metric. A trend in recent work is for *intelligent* machines to run an optimization scheme in the loop with other adaptive components. An example is in human-in-the-loop optimization of assistive devices, where the optimizer tries to improve the human's gait or metabolic usage by actuating their limbs.

A challenge in this line of work is that the machine does not take the human's adaptation into account, effectively assuming that human's response is stationary (in a stochastic or variational sense). The strategic component of these systems are un-modeled, which can lead to sub-optimal performance (e.g. slow convergence or cycling). By explicitly modeling the human-machine system as a *game*, we are able to predict behavior of the learning agents as they interact with one other to seek equilibrium. Specifically, we can predict stability or instability of stationary points in human-machine interaction by spectral analysis of the learning dynamics.

II. Mathematical model

Consider a game played by two strategic agents – a human and a machine – that independently choose their component of a continuous decisions variable $(H, M) \in \mathcal{H} \times \mathcal{M}$. A reasonable adaptation strategy (c.f. [1]–[3]) posits that agents iteratively continuously descent *cost* functions c_H, c_M,

$$\frac{\partial}{\partial t} H = -\alpha \frac{\partial c_H}{\partial H}(H, M), \; \frac{\partial}{\partial t} M = -\beta \frac{\partial c_M}{\partial M}(H, M), \qquad (1)$$

wherein *learning rates* α, β determine small adjustments made to continually decrease c_M and c_H. Despite the simplicity of this descent strategy, the game dynamics in (1), can be complex due to interactions between players. For instance, and in contrast to gradient descent with a single cost function, player costs are *not* guaranteed to decrease at each step. Recent work by ourselves and others [4] provides guarantees on convergence and performance for (1).

III. Experimental setup

We are designing an experiment to test whether game dynamics can predict the behavior of human-machine systems. In particular, we demonstrate that for a fixed game (c_H, c_M), which is chosen to be quadratic, the size of the learning parameter β can elicit categorically different behavior as predicted by the various outcomes of simultaneous (Nash) and sequential (Stackelberg) play.

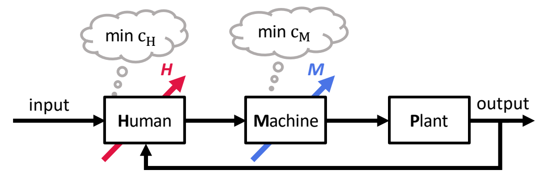

Fig. 1. Human-machine interaction is a closed-loop system. In this diagram, there are two agents: the human and the machine. They depend on each other when controlling the plant. The arrows behind the blocks signify adaptation.

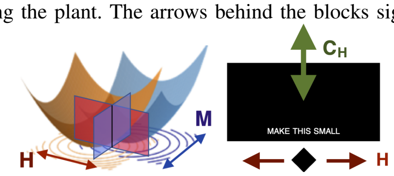

Fig. 2. Human and machine coadaptation can be modeled as a mathematical game. We visualize the quadratic cost landscape of two agents. To solve this game, agents seek stationary points (H^*, M^*) such that they are each minimizers of their own costs, i.e. a Nash or Stackelberg equilibrium.

In the experiment, each player chooses a scalar-valued action $H, M \in \mathbb{R}$ at a frequency of 40Hz for one minute. The interaction is determined by the game vector field: the machine does gradient descent with a constant step size; the human is instructed to move its cursor horizontally to decrease the value of its cost $c_H(H, M)$, which is prescribed to the user through an interface (Figure 2).

To answer the question posed in the title, we do not need to model exactly what learning strategy the human employs. Instead, we can observe the steady state of the system, whether it is stationary, oscillating, or divergent. We strongly believe that game dynamics will serve as an crucial paradigm for modeling intelligent machines that are in-the-loop with humans and with society as a whole. We hope to get feedback from the community on our experiment design.

References

[1] S. Russell and E. Wefald, "Principles of metareasoning," *Artificial intelligence*, vol. 49, 1991.

[2] S. J. Gershman, E. J. Horvitz, and J. B. Tenenbaum, "Computational rationality: A converging paradigm for intelligence in brains, minds, and machines," *Science*, vol. 349, no. 6245, pp. 273–278, 2015.

[3] C. Papadimitriou and G. Piliouras, "Game dynamics as the meaning of a game," *ACM SIGecom Exchanges*, vol. 16, no. 2, pp. 53–63, 2019.

[4] B. J. Chasnov, D. Calderone, B. Açıkmeşe, S. A. Burden, and L. J. Ratliff, "Stability of gradient learning dynamics in continuous games: Scalar action spaces," in *2020 59th IEEE Conference on Decision and Control (CDC)*, IEEE, 2020, pp. 3543–3548.