Co-adaptation improves performance in a dynamic human-machine interface

Momona Yamagami[1,2], Maneeshika M. Madduri[1], Benjamin J. Chasnov[1], Amber H.Y. Chou[1], Lauren N. Peterson[1,2], Samuel A. Burden*[1]

[1]Department of Electrical & Computer Engineering
University of Washington, Seattle, WA 98195, USA
[2]Department of Electrical and Computer Engineering
Rice University, Houston, TX, 77005
*corresponding author: sburden@uw.edu

Abstract

Despite the growing prevalence of adaptive systems in daily life, methods for analysis and synthesis of these systems are limited. Here we find theoretical obstacles to creating optimization-based algorithms that co-adapt with people in the presence of dynamic machines. These theoretical limitations motivate us to conduct human subjects experiments with adaptive interfaces, where we find an interface that decreases human effort while improving closed-loop system performance during interaction with a machine that has complex dynamics. Finally, we conduct computational simulations and find a parsimonious model for the human's adaptation strategy in our experiments, providing a hypothesis that can be tested in future studies. Our results highlight major gaps in understanding of co-adaptation in dynamic human-machine interfaces that warrant further investigation. New theory and algorithms are needed to ensure interfaces are safe, accessible, and useful.

With practice and effort, people can learn to control dynamic systems like vehicles,[1–3] tele- and co-robots,[4–6] prostheses,[7–9] exoskeletons,[10–12] or brain-computer interfaces.[13–15] To facilitate and shape this learning, it is tempting to inject intelligence into the *interface* between the human and dynamic machine.[16] But since people continually adapt to their sensorimotor context,[17,18] introducing an adaptive interface into this closed-loop interaction creates a *two learner problem*[19,20] wherein the human and interface *co-adapt*.[21] Understanding how to analyze and synthesize these systems is critically important in current and emerging applications including driver assistance,[22] surgical robotics,[23] rehabilitative robotics,[24] active prosthetics,[8] and neural interfaces (both invasive[13] and non-invasive[25]).

Motivated by the optimal feedback hypothesis for human motor control,[26–29] we model the interaction between a human, adaptive interface, and dynamic machine using a robust control framework.[30] In particular, we assume the human and interface both solve optimal control problems to determine their behavior. By analyzing the equilibrium behavior of such systems, we find a fundamental theoretical limitation neglected in prior work,[31,32] where it was assumed that the human and interface had noise-free observations of the machine's state vector, which we regard as unrealistic in real-world applications. For example, in the context of neuroprosthetics, where a neural interface seeks to help a person control a dynamic machine, intrinsic stochasticity of neural signals[33] injects noise into the system, precluding access to full information. This finding leads us to conduct experimental and simulation studies to discover and model how people behave when interacting with adaptive interfaces and dynamic machines under the realistic condition where measurements are noisy.

We consider systems where the machine M has dynamics, that is, where the dimension of the machine's state vector – the system *order* [34, Sec. 2.2] – is greater than or equal to 1; the zero-dimensional case was studied in prior work.[35–37] In our theoretical analysis, we allow arbitrarily-large system orders. In our experiments and simulations, we test interfaces that are 0th-, 1st-, and 2nd-order with a specific 2nd-order dynamic machine that is fundamentally challenging to control[2] – a nonminimum phase 2nd-order system. The interface orders we test correspond to the position-, velocity-, and acceleration-based dynamics routinely encountered in daily life.[1]

A model of co-adaptation in dynamic human-machine interfaces

We model dynamic human-machine interfaces (HMI) using the block diagrams in Fig. 1, where: H represents the *human* "in-the-loop", M the *machine* that is being controlled, and I the *interface* we seek to synthesize. These diagrams specify the flow of information, with signals illustrated by *arrows* and transformations of signals illustrated by *blocks*. When the blocks H, M, and I are linear time-invariant (LTI) transformations [38, Lec. 3], these diagrams are not solely conceptual – they provide precise mathematical specifications of the closed-loop transformation from input disturbance w to output error z. Although humans are generally nonlinear, they can behave remarkably linearly when interacting with finite-order LTI machine-and-interface dynamics.[1] These observations motivate us in what follows to focus on the finite-order LTI case where there exists a comprehensive toolkit for analysis and synthesis of feedback systems.[30,34,38]

The field of *robust control theory*[30] provides methods for synthesizing the interface I in Fig. 1 (conventionally termed the *controller*) to optimize a performance criterion with respect to fixed models of human H and machine M (and, optionally, uncertainty in the models). The performance criteria of interest in robust control are *induced norms* [30, Ch. 4] that quantify how much signal power is transferred from input disturbance w to output error z. In this robust control paradigm, the interface I is synthesized by solving an optimization problem

$$\min_I \|H/M/I\|_{\mathcal{I}} \tag{1}$$

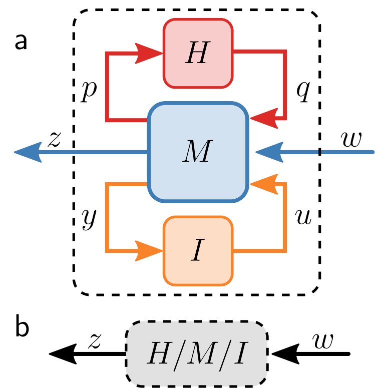

Figure 1: *Block diagram models for human-machine interfaces (HMI).* (a) This diagram specifies that the human H transforms signal p to signal q, i.e. the human observes output p from machine M and provides control input q to the machine. Similarly, the interface I transforms machine output y to control input u, and the machine M transforms both of the control inputs q, u and an external input w to produce outputs p, z, y. The external input w can contain a *disturbance* to reject (e.g. measurement or process noise) or a *reference* to track (e.g. a trajectory or stationary point). (b) This diagram shows a simplification of (a) that obscures details about the interconnection between H, M, and I, instead denoting the transformation $H/M/I$ resulting from this interconnection. The block $H/M/I$ is defined so that (a,b) both specify the same transformation from w to z.

where the induced norm $\|\cdot\|_{\mathcal{I}}$ encodes what components of z the interface seeks to make as small as possible. Under appropriate restrictions on the choice of norm $\|\cdot\|_{\mathcal{I}}$, models of human H and machine M, and statistics of input w and output z, a solution I^* to the optimization problem in (1) exists[30] and can be computed using efficient numerical algorithms.[39] As an example, the well-known *linear-quadratic Gaussian* (LQG) regulator [38, Lec. 24] is obtained by solving (1) when: the disturbance w is Gaussian; the error z consists of (linear transformations of) the state of the machine x and the control input u; the human-machine transformation H/M defined in Fig. 2 is stabilizable and observable [38, Lec. 14, 15]; and $\|\cdot\|_{\mathcal{I}}$ is the induced 2-norm.

If the statistics of the human H and machine M in Fig. 1 are stationary regardless of the implemented interface I (e.g. if H and M are given as fixed transformations or distributions of transformations), then the preceding paragraph describes a remarkably flexible framework for synthesizing an optimal interface I^*.[30,40] However, although it may be reasonable to assume or ensure machine dynamics are stationary, ample evidence suggests that the human will naturally adapt to any perceived change in the interface, and moreover that this adaptation will not be random – rather, the human's transformation will be strongly influenced by the interface.[1,2,25,41] Therefore if the interface I^* is synthesized by solving (1) with respect to an initial guess or estimate of the human's transformation H, it is reasonable to expect the human will adapt its transformation from H to $\widetilde{H}$ when the interface changes from I to I^*. Unfortunately, the synthesized interface I^* is not optimal with respect to the adapted human transformation $\widetilde{H}$.

In fact, implementing I^* in-the-loop with $\widetilde{H}$ could yield arbitrarily bad performance.[42]

The preceding observations motivate the study of interfaces that *co-adapt* with the human and, hence, regard the human H and interface I as two *learners*[19,20] playing a *dynamic game*[43] through their interaction with the machine M. As a starting point for modeling this interaction, prior work suggests the human may play this game by solving their own optimization problem[26–29]

$$\min_{H} \|H/M/I\|_{\mathcal{H}} \tag{2}$$

where the norm $\|\cdot\|_{\mathcal{H}}$ encodes what components of z the human seeks to make as small as possible. Other than in the special case where the goals of the human and interface are perfectly aligned so that $\|\cdot\|_{\mathcal{I}} = \|\cdot\|_{\mathcal{H}}$, the outcome of the game defined by simultaneously considering the optimization problems in (1) and (2) will generally represent a compromise between the player's conflicting goals. One such outcome considered in prior work on human-machine interfaces[19,31,32,35] is a *Nash equilibrium*[43,44] defined by a pair of transformations H^*, I^* such that H^* minimizes $\|H^*/M/I^*\|_{\mathcal{H}}$ and I^* minimizes $\|H^*/M/I^*\|_{\mathcal{I}}$.

Results

Theory results

This section provides theoretical analysis of the dynamic game defined in the preceding section that is played by a human H and interface I interacting with a machine M, where we assume H, M, and I are linear time-invariant (LTI) transformations. The game is specified by the coupled optimization problems

$$\min_{H} \|H/(M/I)\|_{\mathcal{H}} \tag{3a}$$

$$\min_{I} \|(H/M)/I\|_{\mathcal{I}} \tag{3b}$$

where $\|\cdot\|_{\mathcal{H}}$, $\|\cdot\|_{\mathcal{I}}$ denote the utility functions of the human and interface, respectively, and the closed-loop transformations $H/(M/I) = H/M/I = (H/M)/I$ are defined in Fig. 2. The simplified block diagrams in Fig. 2 are useful to illustrate (i) the combined machine-interface system (M/I) the human H interacts with and (ii) the combined human-machine system (H/M) the interface I interacts with.

Given LTI transformations M and I, the solution H^* of the optimization problem in (3a) is an LTI transformation [30, Thm 14.7] that can be computed using efficient algorithms.[39] With the exception of the special cases considered in prior work[31,32] (termed *full information* and *full control* in[30,45]), the number of state variables in the dynamics of H^* is equal to the sum of the number of state variables in M and I; if we let $\#(T)$ denote the number of state variables in the LTI transformation T, this property can be written $\#(H) = \#(M) + \#(I)$. Similarly, given LTI transformations H and M, the solution I^* of the optimization problem in (3b) is an

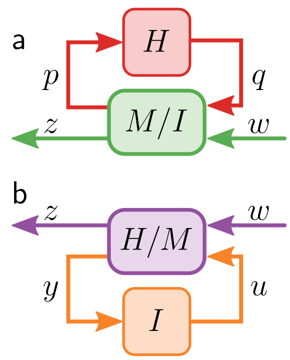

Figure 2: *Block diagrams from human and machine perspectives.* Given LTI transformations H, M, and I, the blocks M/I and H/M are defined so that the diagrams in (a, b) specify the same transformation from w to z as the diagrams in Fig. 1. Mathematically, the U/L operation is defined as the *linear fractional transformation* [30, Ch. 3, 10] between blocks U and L. These simplified diagrams are conceptually useful when reasoning from the individual perspectives of the human H and interface I as they jointly interact with the machine M. Indeed, H interacts with the interconnection M/I between M and I as illustrated in (a), whereas I interacts with the interconnection H/M between H and M as in (b).

LTI transformation, and the number of state variables in I^* equals the sum of the number of state variables in H and M: $\#(I) = \#(H) + \#(M)$.

We will show in the Theorem below that the game in (3) generally has no Nash equilibrium when M has dynamics, i.e., $\#(M) \geq 1$. To see why this may be the case, consider the co-adaptive interaction wherein H and I alternately solve their optimization problems ((3a) and (3b), respectively). Even starting from an interface with $\#(I) = 0$, solving (3a) with respect to given M and I yields solution H^* with $\#(H) = \#(M) \geq 1$, so subsequently solving (3b) with respect to given H^* and M yields solution I^* with $\#(I) = \#(H) + \#(M) = 2\,\#(M) \geq 2$. Iterating this process yields a sequence of H and I with ever-increasing numbers of state variables, preventing the existence of a stationary point for (3) in the sense of Nash.

Theorem: *Suppose M is a linear time-invariant (LTI) transformation with dynamics so that $\#(M) \geq 1$. Then the dynamic game in* (3) *can have a Nash equilibrium H^*, I^* only if the transformations H^*/M and M/I^* are full-information or full-control.*[30,45]

Proof (by contradiction): *Suppose there exists a Nash equilibrium H^*, I^* for* (3)*. If either H^*/M or M/I^* are not full-information or full-control, then [30, Thm. 14.7] implies $\#(H^*) = \#(M)+\#(I^*)$ and $\#(I^*) = \#(H^*)+\#(M)$. Substituting the second equation from the preceding sentence into the first and simplifying yields $0 = 2\,\#(M)$. But since $\#(M) \geq 1$, this equation is a contradiction.*

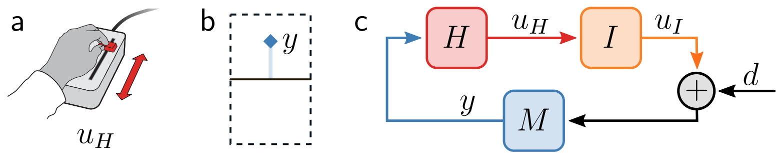

Figure 3: *Experiment design.* (a) Human subjects provide response u_H using a 1-dimensional manual device. (b) The subjects are instructed to change their response to make a cursor on a computer display as close as possible to a reference position in the middle of the screen. (c) The human H, machine M, and interface I are connected in series, with the human viewing output y from the machine M and producing response u_H that is input to the interface I. The interface's response u_I is corrupted by an external disturbance d before being input to the machine.

Remark: *The co-adaptation games studied in prior work[19, 31, 32, 35] consider only the full-information case, where it is assumed that the state of the machine M is observed by both the human H and interface I with no measurement noise. In this case, the game can admit a Nash equilibrium defined in terms of static state feedback transformations for H and I [43, Sec. 6.2.2].*

Example: *The neuroprosthetic example from the introduction illustrates why full-information or full-control assumptions are unrealistic. Indeed, full-information would require both adaptive agents – human and interface – have noise-free measurements of all system states, including those internal to the other adaptive agent. Similarly, full-control would imply both agents can directly influence all system states; although it might be possible in principle to give the brain full control over the machine and interface, it seems implausible (and perhaps undesirable) for the interface to have full control over the human's neural state. Either assumption seems inconsistent with our understanding of neural dynamics.[33]*

Experiment results

We tested the effect of co-adaptation on HMI performance in the presence of the machine with dynamics in (6), where neither the human or the interface have full-information or full-control. Participants in our human subjects experiment completed a continuous disturbance-rejection task using a one-dimensional manual input device[1,3] as shown in Fig. 3. We assessed performance using three metrics: how well the HMI rejected the input disturbance, quantified using the transformation norm $\|H/M/I\|^2$; how much effort the human exerted, quantified as $\|H\|^2$; and how much effort the interface exerted, quantified as $\|I\|^2$. In each case, we used the induced 2-norm of the transformation. The interface adapted by minimizing a linear combination of task performance and interface effort,

$$c_I(H, I) = \|H/M/I\|^2 + \lambda_I \|I\|^2, \tag{4}$$

with $\lambda_I = 10^{-4}$.

We recruited eleven participants, all of whom were daily computer users. Participants completed the co-adaptive disturbance-rejection task over a sequence of 21 trials in each of three conditions presented in random order: i) 0th-order interface; ii) 1st-order interface; iii) 2nd-order interface (Figure 4a). Each condition started with a randomized interface initialization, and the interface was adapted every three trials by estimating a model of the human transformation $\widetilde{H}$ and minimizing $c_I(\widetilde{H}, I)$ with respect to I. Between each condition, participants completed three *baseline* trials where the interface was set to a constant unity gain (passthrough, $I = 1$). Participants were asked to keep the randomly disturbed cursor as close to the center of the screen as possible (see Fig. 3b).

We found that the co-adaptive 1st-order interface improved both the task performance and human effort metrics: $\|H/M/I\|^2$ and $\|H\|^2$ decreased significantly relative to baseline (Figure 4b,c *top*; $^*P < 0.05$; Wilcoxon signed-rank test). The interface effort *increased* significantly relative to baseline (Figure 4e *top*; $^*P < 0.05$; Wilcoxon signed-rank test), but this increase was compensated by the decrease in $\|H/M/I\|^2$ so that the interface's cost in (4) decreased significantly (Figure 4d; $P = 0.04$; Wilcoxon signed-rank test). We also found that the task performance after co-adaptation with the 1st-order interface was indistinguishable from the baseline performance with the 0th-order interface (Figure S1 top row; $P = 0.23$; Wilcoxon signed-rank test) but significantly better than the 2nd-order interface (Figure S1 top row; $^*P = 0.03$; Wilcoxon signed-rank test).

Observing the spectral density plots more carefully, we found there was no statistically significant difference between co-adaptation and baseline at each individual tested frequency for task performance (Figure 4b, bottom; $P > 0.05$; Wilcoxon signed-rank test). However, the median of the baseline was higher after co-adaptation at low frequencies below the *crossover frequency* 0.25 Hz where prior work has shown humans adapt significantly.[1,3] Similarly, human effort was lower after co-adaptation compared to baseline at lower frequencies (below crossover) but not at higher frequencies (Figure 4c; $^*P < 0.05$; Wilcoxon signed-rank test). The interface had higher effort at all stimulated frequencies with co-adaptation compared to baseline (Figure 4e; $^*P < 0.05$; Wilcoxon signed-rank test). These findings are consistent with prior work that demonstrates the human operator's ability to reject disturbances falls off rapidly for frequencies above crossover,[1,3] whereas adaptive interfaces do not have such frequency-based restrictions. Together, these results show that the co-adaptive 1st-order interface improved disturbance-rejection task performance while decreasing human effort and increasing interface effort.

Simulation results

Based on the preceding theory and experiment results, we conducted a simulation study to obtain a parsimonious computational model of the human and 1st-order interface co-adaptation. Our goal was to determine human parameters that approximated our experimental results. Since our theory results show that there can be no Nash equilibrium when the order of the

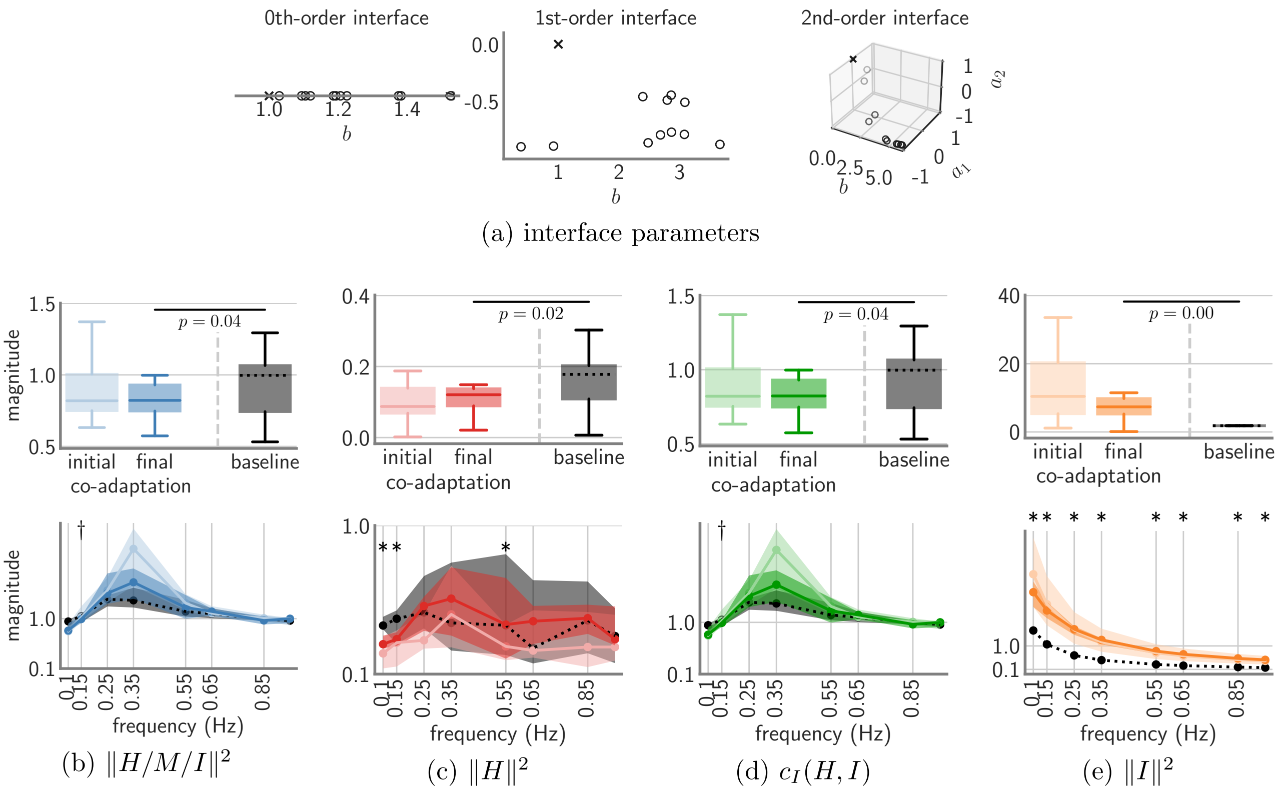

Figure 4: *Experiment results (N=11 participants).* **(a)** Baseline and adapted interface parameters. The 0th-, 1st-, and 2nd-order adaptive interfaces converged to a range of different parameters (*left*, *middle*, and *right*, respectively). Baseline interface ($I = 1$) shown by cross (x); final adapted interfaces shown by open circles (∘) for each participant. **(b-d)** Distributions of performance metrics for baseline and adapted interfaces shown using box plots (*top*; 0th, 25th, 50th, 75th, and 100th percentiles) and spectral density plots (*bottom*; median and interquartile): task performance $\|H/M/I\|^2$ in **(b)**; human effort $\|H\|^2$ in **(c)**; interface cost $c_I(H, I)$ in **(d)**; interface effort $\|I\|^2$ in **(e)**. Statistically significant differences between baseline and final adapted interfaces shown with horizontal lines in box plots and with asterisk (*) spectral density plots ($P < 0.05$, Wilcoxon signed-rank test); significant differences between initial and final adapted interfaces shown with dagger (†) in spectral density plots ($P < 0.05$, Wilcoxon signed-rank test).

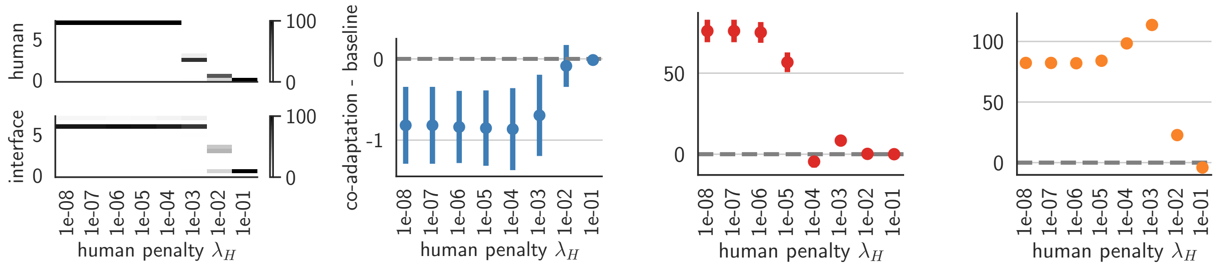

(a) transformation gain b (b) $\|H/M/I\|^2$ (c) $\|H\|^2$ (d) $\|I\|^2$

Figure 5: *Simulation results (N=100 random initializations).* **(a)** Gain parameter (b) from co-adaptation simulation outcomes for human (*top*) and interface (*bottom*) transformations over human penalty parameter λ_H; human model H is 2nd-order, interface I is 1st-order. Intensity denotes percentage of simulation outcomes. **(b-d)** Distributions of differences between performance metrics for baseline and adapted HMI shown using mean +/- standard deviation: task performance $\|H/M/I\|^2$ in **(b)**; human effort $\|H\|^2$ in **(c)**; interface effort $\|I\|^2$ in **(d)**. Baseline indicated with grey dashed line. A negative value indicates that the simulated co-adaptation magnitude was lower than the baseline magnitude. A positive value indicates that the simulated co-adaptation magnitude was higher than the baseline magnitude.

human and interface dynamics are unbounded, we only simulate low-order dynamics for both. Following prior work,[35,36] we assumed that the human adapted their parameters by minimizing a linear combination of task performance and human effort:

$$c_H(H, I) = \|H/M/I\|^2 + \lambda_H \|H\|^2, \tag{5}$$

where we again used induced 2-norms. We simulated the human and interface co-adaptation for human penalty parameters λ_H ranging from 10^{-8} to 10^{-1}. The 2nd-order machine M had the same parameterization as in the experiment, (6). We only simulated co-adaptation with the 1st-order interface, since that was the only experimental condition that resulted in significant changes in task performance.

To simulate human and interface co-adaptation, we randomly initialized the human and 1st-order interface within the parameter ranges tested in the experiments. Then, we synthesized the optimal human model H^* for the optimization problem in (5). Next, we held the human model constant and synthesized the optimal interface I^* for the optimization problem in (4). We continued alternating optimizing for the human and interface until parameter convergence. We repeated the random parameter initialization and alternating optimization 100 times to determine whether initialization location affected the final parameters after co-adaptation. We additionally performed the same optimization but solely for the human model, holding the interface constant at $\widehat{I} = 1$ to obtain how the human model adapted for a baseline interface.

A 2nd-order human model with a penalty in the middle of the range we tested ($\lambda_H = 10^{-4}$) yields simulation outcomes that qualitatively correspond to our experimental results, where we found co-adaptation improved task performance, decreased human effort, and increased interface effort relative to baseline (Figure 5b-d). In contrast, simulation outcomes with 0th- and 1st-order human models were inconsistent with one or more of our experimental results across all tested human penalty parameters (Figure S3). With a 2nd-order human model, both the interface and human had multiple Nash equilibria across all human penalties tested (Figure 5a). Lower human penalty resulted in higher human and interface gains, whereas higher human penalty resulted in lower human and interface gains.

Discussion

Our experiment results demonstrate potential advantages – and potential drawbacks – of deploying adaptive interfaces in-the-loop with humans and dynamic machines. We found co-adaptation with a 1st-order interface significantly improved task performance while lowering human effort and decreasing interface cost relative to baseline (Figure 4b,c,d). The decrease in interface cost was obtained despite an *increase* in interface effort (Figure 4e), representing a compromise between the interface's opposing goals of improving task performance with minimum effort, (4). Game-theory methods for co-adaptation that explicitly consider strategic interests of two intelligent decision-making agents may enable the interface designer to systematically explore this tradeoff.[37]

The performance improvements in our experiments were achieved using interfaces that were tailored to the individual (Figure 4a); it is unclear whether similar improvements would have obtained by optimizing a single interface for the entire population of human subjects. However, the improvements we observed were by no means guaranteed, as co-adaptation between intelligent decision-making agents may yield worse outcomes for all or fail to converge entirely.[21,46] Indeed, the fact that we did not observe significant improvements for 0th- and 2nd-order interfaces demonstrates the influence of the adaptive interface's parameterization on outcomes. Additionally, prior results highlight the importance of giving people sufficient time to learn.[3,41]

Our simulation results (Figure 5) corroborate the longstanding observation that humans themselves have dynamics.[1–3,41] Coupled with our theory result, that co-adaptation between optimal agents generally has no Nash equilibrium when system orders are unconstrained, this simulation finding implies that humans do not implement the theoretically-optimal feedback controller[26–29] when interacting with machines that have complex dynamics. By restricting the generally ill-posed game in (3) to finite-order human and interface transformations, we found one or more Nash equilibria in our simulations (Figure S2), and one particular parameterization of the human transformation and cost that yielded outcomes consistent with our experimental results. Our simulation framework may prove useful to test hypotheses and inform interface design in future studies.

Finally, our theory result brings to light a fundamental limitation in current understanding of how to analyze and synthesize interfaces that co-adapt in-the-loop with a human and dynamic machine. Importantly, real-world co-adaptive systems will generally be neither full-information or full-control due to noise in sensory-and-motor channels and inability of the interface to directly control all system states – especially those of the human's dynamics. So we claim that our theory result has broad practical implications. As an avenue for future exploration, it is important to note that our results concerned only Nash equilibria, but dynamic games can yield a rich variety of outcomes depending on the information structure, order of play, and strategy employed by players.[37] A new paradigm is needed to ensure safety,[47,48] accessibility,[49,50] and utility[51–53] of co-adaptation between humans and intelligent interfaces.

Conclusion

Our work highlights limitations in prior theory, experiment, and simulation work on co-adaptation between humans and intelligent interfaces interacting in closed loop with dynamic machines in real-world conditions. We demonstrate that a Nash equilibrium does not exist under standard assumptions of optimal adaptation, and explore in experiment and simulation how a human and interface with *bounded rationality*[54] adapt to improve performance in a disturbance-rejection task. Understanding how to analyze and synthesize co-adaptive systems, where both human and algorithmic agents adapt to control dynamic systems like vehicles, robots, and prostheses is critically important in current and emerging applications including driver assistance, rehabilitation robotics, and neuroprosthetics. We contribute new theory, experiment, and simulation results that will inform creation of these systems for real-world deployment.

Materials and methods

Data and analyses to reproduce the results reported here are available in a permanent publicly-accessible repository.[55]

Human subjects

All participants provided informed consent according to the University of Washington, Seattle's Institutional Review Board (IRB #00000909). The goal of this experiment was to determine the effects of human and interface co-adaptation on final interface dynamics, task performance, and human and interface effort. This was a pilot study to determine whether and how the HMI co-adapt and how the co-adaptation affects HMI performance and human and interface effort compared to baseline. Eleven participants were recruited for the study (age: 28 ± 7 years (mean $\pm$ standard deviation); gender: 8 women, 3 men, 1 non-binary (some identified with multiple genders); hand dominance: 11 right-hand dominant). All were daily computer users.

Task

Participants were tasked with controlling a cursor on the screen with a one-dimensional slider, which was built from a $35 \times 12 \times 22$ mm (width$\times$height$\times$depth) rectangular handle attached to a slide potentiometer with a 10 cm extent (Figure 3a). Following prior work,[3,56] unpredictable disturbance signals d were constructed as a sum of sinusoidal signals with the first eight prime multiples of a base frequency of 1/20 Hz ($\Omega = [0.10, 0.15, 0.25, 0.35, 0.55, 0.65, 0.85, 0.95]$ Hz). Each frequency component's magnitude was normalized by the frequency squared to ensure constant signal power, and the phase of each frequency component was randomized in each trial to produce pseudorandom time-domain signals. The disturbances perturbed the cursor in an unpredictable fashion, and participants were asked to keep the cursor as close to the center of the screen as possible. Each trial was 40 seconds after a 5-second ramp-up.

Human response u_H were transformed through a fixed non-minimum phase second-order machine M to produce output y to increase the complexity of the task:[2]

$$\begin{aligned} \ddot{y} + 3.6\dot{y} + 4 &= 2(\dot{u_I} + 2.2u_I) + d, \\ \widehat{M}(s) &= \frac{2(s+2.2)}{s^2 + 3.6s + 4}, \end{aligned} \tag{6}$$

where $\dot{x}$ represents the time-domain signal x differentiated by time t, and $\widehat{M}(s)$ represents the Fourir transform of machine $M(t)$. As the updates to the cursor position y occurred at 60 Hz, the machine was discretized prior to implementation.

Conditions

We tested three interface conditions (0th-, 1st-, and 2nd-order; Table 1). As the updates to the interface response u_I occurred at 60 Hz, discrete dynamics were used to update the interface output from one time point to the next. Each condition started with three trials of a baseline where the human response u_H was unaffected by the interface dynamics ($u_I[t] = u_H[t-1]$, $\widehat{I}_{baseline}(z) = 1$), followed by 21 trials of co-adaptation for each condition (∼30 minutes per condition). The condition order was randomized for each participant. After all three conditions were completed, the participants performed three more baseline trials. All participants were encouraged to take breaks between the 45-second trials, and participants were asked to take at least a one-minute break after each condition.

Table 1: *0th-, 1st-, and 2nd-order interfaces tested in experiments.*

	time domain	**frequency domain**
0th	$u_I[t] = bu_H[t-1]$	$\widehat{I}_0(z) = b$
1st	$u_I[t] = au_I[t-1] + bu[t-1]$	$\widehat{I}_1(z) = \frac{b}{z-a}$
2nd	$u_I[t] = (a_1 + a_2)u_I[t-1]$ $-a_1a_2u_I[t-2] + bu[t-1]$	$\widehat{I}_2(z) = \frac{b}{(z-a_1)(z-a_2)}$

Signal processing

Prior work demonstrated that when humans are tasked with tracking references r and rejecting additive disturbances d through a linear time-invariant (LTI) [34, Ch. 3, pg. 4] system M, humans behave approximately like LTI transformations for a range of reference and disturbance signals.[1,3,56] As such, we can analyze our system in Figure 3c using the frequency-domain representations [57, Ch. 5] of signals and LTI systems; we will adorn signal x and transformation T with a "hat" $\widehat{\cdot}$ to denote the Fourier transform $\widehat{x}$, $\widehat{T}$. Therefore, for a given prescribed and measured signals and transformations $\widehat{d}, \widehat{y}, \widehat{I}, \widehat{M}$, we can apply block diagram algebra [34, Sec. 2.2] to transcribe Figure 3c into equations that can then be manipulated to express the empirical and prescribed transfer functions $\widehat{T}_{u_H d} = \frac{\widehat{u}_H}{\widehat{d}}$ as a function of the unknown human transfer function $\widehat{H}(\omega)$:

$$\widehat{u}_H(\omega) = \underbrace{\frac{-\widehat{H}(\omega)\widehat{M}(\omega)}{1+\widehat{H}(\omega)\widehat{M}(\omega)\widehat{I}(\omega)}}_{\widehat{T}_{u_H d}(\omega)}\widehat{d}(\omega), \tag{7}$$

where $\omega \in \Omega$. We can then estimate the human's controller $\widehat{H}(\omega)$ at specific stimulus frequencies ω as:

$$\widehat{H}(\omega) = -\widehat{M}^{-1}(\omega)\frac{\widehat{T}_{u_H d}(\omega)}{1+\widehat{I}(\omega)\widehat{T}_{u_H d}(\omega)} \tag{8}$$

We can additionally apply block diagram algebra to obtain the human- and interface-controlled cursor position $\widehat{y}$ as a function of prescribed and measured signals and transfer functions:

$$\widehat{y}(\omega) = \frac{\widehat{M}(\omega)}{1+\widehat{H}(\omega)\widehat{M}(\omega)\widehat{I}(\omega)}\widehat{d}(\omega) \tag{9}$$

Lastly, we define our performance metric and corresponding cost function c_I in (4) for our specific task at hand:

$$\|H/M/I\| = \left\|\frac{\widehat{y}}{\widehat{d}}\right\| \tag{10}$$

where $\|\cdot\|$ represents the induced 2-norm.

Interface adaptation

The interface update occurred every three trials for a total of 21 trials for each condition. This was to ensure that participants had sufficient exposure to the new interface and adapt, and was based on results obtained from pilot studies (not shown). Interface $\widehat{I}(\omega)$ parameters were restricted such that $0.2 < b < 7$, $-0.95 < a, a_1, a_2 < 0.7$ and and initially randomly assigned within those ranges. We initially restricted the poles to have any magnitude less than 1 to ensure a stable interface, but found during pilot studies that poles smaller than -0.95 or larger than 0.7 resulted in an uncontrollable interface. In addition, we initially did not restrict the 2nd-order interface poles to be solely real numbers, but found during pilot studies that participants were only converging to real values and so chose to only search over real poles.

The interface updates occurred in four steps. First, three trials of the disturbance rejection task were completed by the participant. Next, the human model $\widehat{H}$ in (8) was estimated by solely using the data from the last two trials. Subsequently, a grid search over the ranges of a, b were conducted to minimize the cost c_I in (4) with $\lambda_I = 10^{-4}$, and the resulting minimizing interface $\widehat{I}^*$ was noted. The grid search was initialized with 100 equidistant points between the ranges noted above. Lastly, to ensure gradual changes between interfaces from trial to trial and slower changes with increasing trial numbers, *Smooth Batch*[13] was implemented. The subsequent interface $\widehat{I}^+$ was defined as a weighted combination of the prior interface $\widehat{I}^-$ and the computed optimal interface $\widehat{I}^*$:

$$\widehat{I}^+ = \alpha \widehat{I}^- + (1-\alpha)\widehat{I}^*. \tag{11}$$

The parameter α was used to adjust the weighting of the prior interface and computed optimal interface and linearly increased from 0 (i.e., subsequent interface is solely the optimal interface) to 1 (i.e., the subsequent interface is solely the prior interface) as the number of trials increased from 0 to 21 trials. This ensured that the interface would update rapidly initially, and then more slowly as the number of trials increased.

Statistical analysis

Our primary outcomes of interest were performance differences between the final co-adaptation and baseline and between the initial and final co-adaptation. To determine whether the differences were statistically significant, we computed the average magnitude of each performance metric of interest (task performance: $\|H/M/I\|^2$, (10); human effort: $\|H\|^2$, (8); interface effort: $\|I\|^2$, (3b); cost: $c_I(\widehat{H}, \widehat{I})$, (4)) with the Wilcoxon signed-rank test [58, Sec. 5.7]. The Wilcoxon signed-rank test was chosen because the residuals of our dataset was not normally distributed ($P > 0.05$; Shapiro Wilk test). Due to an experimenter error, only 10 out of the 11 participants were analyzed for the 2nd-order interface.

Simulation methods

The goal of the simulation was to develop a predictive model of human co-adaptation with a 1st-order interface and establish simulation parameters that approximate our experimental results. Towards this goal, we tested co-adaptation of three human model parameterizations (0th-, 1st-, and 2nd-order) with a 1st-order adaptive interface and 2nd-order fixed machine for various human penalty terms λ_H ranging from 10^{-8} to 10^{-1}. The 2nd-order fixed machine M had the same parameterization as in the experiment, (6), and was defined as a non-minimum phase 2nd-order dynamical system.

Co-adaptation and baseline simulations

We sequentially optimized the human and interface for their respective cost functions. We first randomly initialized the human and interface within the parameter bounds tested experimentally (Table 1; $0.2 < b < 7, -0.95 < a, a_1, a_2 < 0.7$). Next, we synthesized the optimal human model H^* by minimizing the human cost $c_H(\widehat{H}, \widehat{I})$ in (5). After, we held the human model constant and synthesized the optimal interface I^* by minimizing the interface cost $c_I(\widehat{H}, \widehat{I})$ in (4). We used the same interface penalty term used in the experiment ($\lambda_I = 10^{-4}$). We repeated the alternating optimization of the human and interface parameters until convergence. To determine how parameter initialization affects the final human and interface parameters after co-adaptation, we repeated the randomized initialization and convergence 100 times. We additionally performed the same optimization but solely for the human model, holding the interface constant at $\widehat{I} = 1$ to obtain how the human model adapted for a baseline interface.

Simulation analysis

Using the final human and interface parameters for each of the 100 initializations and tested human penalty λ_H, we computed performance metric (task performance: $\|H/M/I\|^2$, (10); human effort: $\|H\|^2$, (8); interface effort: $\|I\|^2$, (3b)) for baseline and co-adaptation. We compared the spread between the co-adaptation and baseline performance by assuming that the two distributions were normally distributed and taking the difference. For each performance metric, we determined the range of human penalty λ_H that fit our experimental results.

Acknowledgements

This work was funded by a American Society for Engineering Education National Defense Science and Engineering Graduate Fellowship (NDSEG) to MMM and National Science Foundation (NSF) awards #2045014,2124608 to SAB.

Author contributions

M.Y, M.M.M., B.J.C., and S.A.B. conceived the experiments, A.H.Y.C. and L.N.P. conducted the experiments, M.Y., M.M.M., and S.A.B. analyzed the results, and M.Y., M.M.M, and S.A.B. wrote the paper. All authors reviewed the manuscript.

Data Availability

Data and analyses to reproduce the results reported here are available in a permanent publicly-accessible repository.[55]

References

[1] Duane T McRuer and Henry R Jex. A review of Quasi–Linear pilot models. *IEEE Transactions on Human Factors in Electronics*, HFE–8(3):231–249, 1967.

[2] Xingye Zhang, T Michael Seigler, and Jesse B Hoagg. The impact of nonminimum-phase zeros on human-in-the-loop control systems. *IEEE Transactions on Cybernetics*, 52(6):5098–5112, June 2022.

[3] Momona Yamagami, Lauren N Peterson, Darrin Howell, Eatai Roth, and Samuel A Burden. Effect of handedness on learned controllers and sensorimotor noise during trajectory-tracking. *IEEE Transactions on Cybernetics*, 53(4):2039–2050, 2023.

[4] Stefanos Nikolaidis, Swaprava Nath, Ariel D Procaccia, and Siddhartha Srinivasa. Game-theoretic modeling of human adaptation in human-robot collaboration. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 323–331, 2017.

[5] Rahul B Warrier and Santosh Devasia. Iterative learning from novice human demonstrations for output tracking. *IEEE Transactions on Human-Machine Systems*, 46(4):510–521, 2016.

[6] Ajmal Zemmar, Andres M Lozano, and Bradley J Nelson. The rise of robots in surgical environments during COVID-19. *Nature Machine Intelligence*, 2(10):566–572, 2020.

[7] Dalia De Santis. A framework for optimizing co-adaptation in body-machine interfaces. *Frontiers in Neurorobotics*, 15:40, 2021.

[8] Katie Z Zhuang, Nicolas Sommer, Vincent Mendez, Saurav Aryan, Emanuele Formento, Edoardo D'Anna, Fiorenzo Artoni, Francesco Petrini, Giuseppe Granata, Giovanni Cannaviello, et al. Shared human–robot proportional control of a dexterous myoelectric prosthesis. *Nature Machine Intelligence*, 1(9):400–411, 2019.

[9] Vahid Azimi, Tony Shu, Huihua Zhao, Rachel Gehlhar, Dan Simon, and Aaron D Ames. Model-based adaptive control of transfemoral prostheses: Theory, simulation, and experiments. *IEEE Transactions on Systems, Man, and Cybernetics*, 51(2):1174–1191, February 2021.

[10] Juanjuan Zhang, Pieter Fiers, Kirby A Witte, Rachel W Jackson, Katherine L Poggensee, Christopher G Atkeson, and Steven H Collins. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science*, 356(6344):1280–1284, June 2017.

[11] Wyatt Felt, Jessica C Selinger, J Maxwell Donelan, and C David Remy. "Body-In-The-Loop": Optimizing device parameters using measures of instantaneous energetic cost. *PloS One*, 10(8):e0135342, 2015.

[12] Patrick Slade, Mykel J Kochenderfer, Scott L Delp, and Steven H Collins. Personalizing exoskeleton assistance while walking in the real world. *Nature*, 610(7931):277–282, October 2022.

[13] Amy L Orsborn, Siddharth Dangi, Helene G Moorman, and Jose M Carmena. Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):468–477, 2012.

[14] Francis R Willett, Donald T Avansino, Leigh R Hochberg, Jaimie M Henderson, and Krishna V Shenoy. High-performance brain-to-text communication via handwriting. *Nature*, 593(7858):249–254, 2021.

[15] Musa Mahmood, Deogratias Mzurikwao, Yun-Soung Kim, Yongkuk Lee, Saswat Mishra, Robert Herbert, Audrey Duarte, Chee Siang Ang, and Woon-Hong Yeo. Fully portable and wireless universal brain–machine interfaces enabled by flexible scalp electronics and deep learning algorithm. *Nature Machine Intelligence*, 1(9):412–422, 2019.

[16] Jenifer Miehlbradt, Alexandre Cherpillod, Stefano Mintchev, Martina Coscia, Fiorenzo Artoni, Dario Floreano, and Silvestro Micera. Data-driven body-machine interface for the accurate control of drones. *Proceedings of the National Academy of Sciences of the United States of America*, 115(31):7913–7918, July 2018.

[17] Jordan A Taylor, John W Krakauer, and Richard B Ivry. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8):3023–3032, 2014.

[18] James B Heald, Máté Lengyel, and Daniel M Wolpert. Contextual inference underlies the learning of sensorimotor repertoires. *Nature*, 600(7889):489–493, December 2021.

[19] Jan Saputra Müller, Carmen Vidaurre, Martijn Schreuder, Frank C Meinecke, Paul Von Bünau, and Klaus-Robert Müller. A mathematical model for the two-learners problem. *Journal of Neural Engineering*, 14(3):036005, 2017.

[20] S Perdikis and J d. R. Millan. Brain-Machine interfaces: A tale of two learners. *IEEE Systems, Man, and Cybernetics Magazine*, 6(3):12–19, July 2020.

[21] Maneeshika M Madduri, Samuel A Burden, and Amy L Orsborn. Biosignal-based co-adaptive user-machine interfaces for motor control. *Current Opinion in Biomedical Engineering (COBME)*, 27:100462, April 2023.

[22] Klaus Bengler, Klaus Dietmayer, Berthold Farber, Markus Maurer, Christoph Stiller, and Hermann Winner. Three decades of driver assistance systems: Review and future perspectives. *IEEE Intelligent Transportation Systems Magazine*, 6(4):6–22, 2014.

[23] Paula Gomes. Surgical robotics: Reviewing the past, analysing the present, imagining the future. *Robotics and Computer-Integrated Manufacturing*, 27(2):261–266, 2011.

[24] Laura Marchal-Crespo and David J Reinkensmeyer. Review of control strategies for robotic movement training after neurologic injury. *Journal of Neuroengineering and Rehabilitation*, 6(1):1–15, 2009.

[25] Momona Yamagami, Katherine M Steele, and Samuel A Burden. Decoding intent with control theory: Comparing muscle versus manual interface performance. In *ACM Conference on Human Factors in Computing Systems (CHI)*, pages 1–12, April 2020.

[26] Jörn Diedrichsen, Reza Shadmehr, and Richard B Ivry. The coordination of movement: Optimal feedback control and beyond. *Trends in Cognitive Sciences*, 14(1):31–39, 2010.

[27] E Todorov and M I Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.

[28] Jessica C Selinger, Jeremy D Wong, Surabhi N Simha, and J Maxwell Donelan. How humans initiate energy optimization and converge on their optimal gaits. *The Journal of Experimental Biology*, 222(Pt 19), October 2019.

[29] Jeremy L Emken, Raul Benitez, Athanasios Sideris, James E Bobrow, and David J Reinkensmeyer. Motor adaptation as a greedy optimization of error and effort. *Journal of Neurophysiology*, 97(6):3997–4006, June 2007.

[30] Kemin Zhou, John Comstock Doyle, and Keith Glover. *Robust and Optimal Control*, volume 40. Prentice Hall, 1996.

[31] Josh S Merel, Roy Fox, Tony Jebara, and Liam Paninski. A multi-agent control framework for co-adaptation in brain-computer interfaces. *Advances in Neural Information Processing Systems (NeuroIPS)*, 26, 2013.

[32] Yanan Li, Gerolamo Carboni, Franck Gonzalez, Domenico Campolo, and Etienne Burdet. Differential game theory for versatile physical human–robot interaction. *Nature Machine Intelligence*, 1(1):36–43, 2019.

[33] Peter Dayan and Laurence F Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT press, 2005.

[34] Karl Johan Aström and Richard M Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton university press, 2010.

[35] Maneeshika M Madduri, Samuel A Burden, and Amy L Orsborn. A game-theoretic model for co-adaptive brain-machine interfaces. In *2021 10th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 327–330. IEEE, 2021.

[36] Maneeshika M Madduri, Momona Yamagami, Augusto X T Millevolte, Si Jia Li, Sasha N Burckhardt, Samuel A Burden, and Amy L Orsborn. Co-adaptive myoelectric interface for continuous control. In *IFAC Workshop on Cyber-Physical-Human Systems (CPHS)*, volume 55(41), pages 95–100, January 2022.

[37] Benjamin J Chasnov, Lillian J Ratliff, and Samuel A Burden. Human adaptation to adaptive machines converges to game-theoretic equilibria. *Submitted; under review.*, May 2023. http://arxiv.org/abs/2305.01124.

[38] João P Hespanha. *Linear Systems Theory*. Princeton University Press, 2009.

[39] Peter Benner, Volker Mehrmann, Vasile Sima, Sabine Van Huffel, and Andras Varga. SLICOT—A subroutine library in systems and control theory. In Biswa Nath Datta, editor, *Applied and Computational Control, Signals, and Circuits*, pages 499–539. Birkhäuser Boston, 1999.

[40] James Anderson, John C Doyle, Steven H Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, January 2019.

[41] Xingye Zhang, Shaoqian Wang, Jesse B Hoagg, and T Michael Seigler. The roles of feedback and feedforward as humans learn to control unknown dynamic systems. *IEEE Transactions on Cybernetics*, 48(2):543–555, February 2018.

[42] John C Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23(4):756–757, 1978.

[43] Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. Number 23 in Classics in applied mathematics. SIAM, Philadelphia, 2nd ed edition, 1999.

[44] J F Nash. Equilibrium points in N-Person games. *Proceedings of the National Academy of Sciences (PNAS)*, 36(1):48–49, January 1950.

[45] John C Doyle, Keith Glover, Pramod P Khargonekar, and Bruce A Francis. State-space solutions to standard H_2 and H_∞ control problems. *IEEE Transactions on Automatic Control*, 34(8):831–847, 1989.

[46] Benjamin Chasnov, Lillian Ratliff, Eric Mazumdar, and Samuel Burden. Convergence analysis of Gradient-Based learning in continuous games. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 115 of *Proceedings of Machine Learning Research*, pages 935–944. PMLR, 2020.

[47] Sanjit A Seshia, Dorsa Sadigh, and S Shankar Sastry. Toward verified artificial intelligence. *Communications of the ACM*, 65(7):46–55, June 2022.

[48] Christian Pek, Stefanie Manzinger, Markus Koschi, and Matthias Althoff. Using online verification to prevent autonomous vehicles from causing accidents. *Nature Machine Intelligence*, 2(9):518–528, September 2020.

[49] Krzysztof Z Gajos, Amy Hurst, and Leah Findlater. Personalized dynamic accessibility. *Interactions*, 19(2):69–73, March 2012.

[50] Jacob O Wobbrock, Shaun K Kane, Krzysztof Z Gajos, Susumu Harada, and Jon Froehlich. Ability-Based Design: Concept, Principles and Examples. *ACM Transactions on Accessible Computing*, 3(3):1–27, April 2011.

[51] Rob Kling. The organizational context of user-centered software designs. *The Mississippi Quarterly*, 1(4):41–52, 1977.

[52] Don Norman. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books, November 2013.

[53] Jodi Forlizzi. Moving beyond user-centered design. *Interactions*, 25(5):22–23, August 2018.

[54] A Tversky and D Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, September 1974.

[55] Momona Yamagami, Maneeshika M Madduri, Benjamin J Chasnov, Amber H Y Chou, Lauren N Peterson, and Samuel A Burden. Co-adaptation improves performance in dynamic human-machine interfaces [source code]. *CodeOcean*, 2023. http://dx.doi.org/10.24433/CO.7860456.v2.

[56] Momona Yamagami, Darrin Howell, Eatai Roth, and Samuel A Burden. Contributions of feedforward and feedback control in a manual trajectory-tracking task. In *IFAC Conference on Cyber-Physical-Human Systems (CPHS)*, volume 51(34), pages 61–66. Elsevier, 2018.

[57] Charles L Phillips, John M Parr, and Eve A Riskin. *Signals, systems, and transforms*. Prentice Hall, 2003.

[58] W J Conover. *Practical Nonparametric Statistics*. John Wiley & Sons, 1999.

Extended data figures

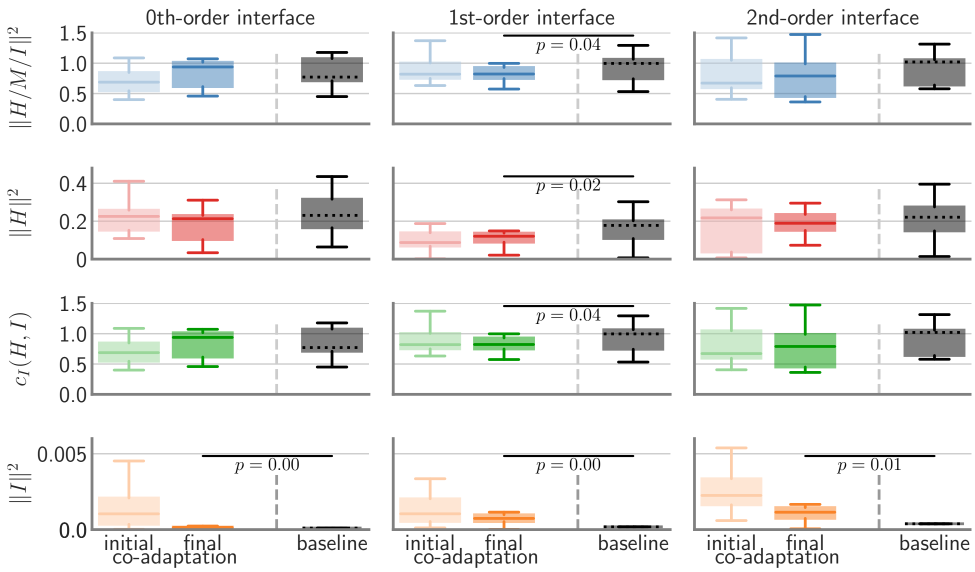

Figure S1: *Experiment results for all interface conditions.* Distributions of performance metrics for baseline and adaptive interfaces shown using box plots (0th, 25th, 50th, 75th, and 100th percentiles) from top to bottom: task performance $\|H/M/I\|^2$; human effort $\|H\|^2$; cost $c_I(H, I) = \|H/M/I\|^2 + \lambda_I \|I\|^2$; interface effort $\|I\|^2$. Each column represents the performance for the 0th-, 1st-, and 2nd-order interface. Statistically significant differences denoted with horizontal lines ($P < 0.05$, Wilcoxon signed-rank test). N=11 participants for 0th- and 1st-order interface; N=10 participants for 2nd-order interface.

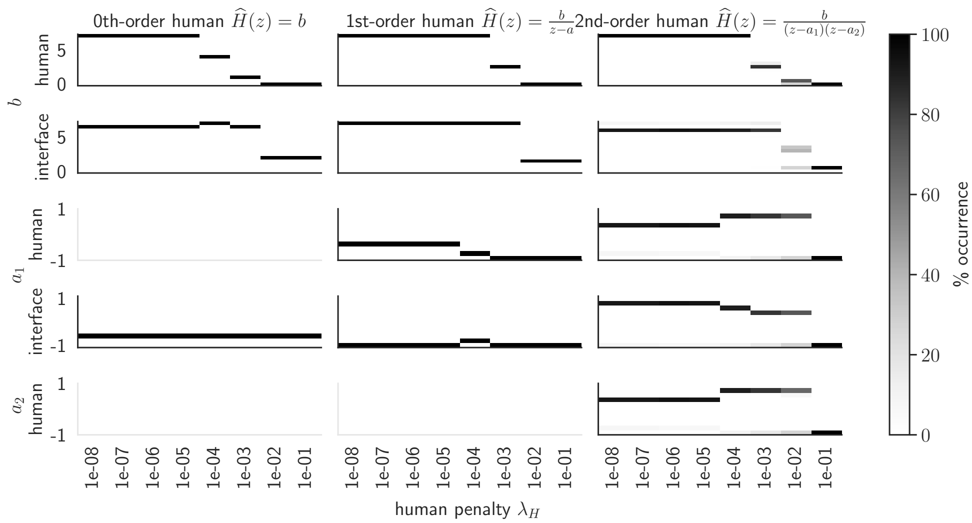

Figure S2: *Simulation results for all interface conditions – Nash equilibria parameters.* Nash equilibria parameters for co-adaptive human (top) and 1st-order interface (bottom) transformations over human penalty parameters λ_H; 0th-, 1st-, and 2nd-order human models $\widehat{H}$ (left to right) are tested. Intensity denotes percentage of simulation outcomes (N=100 random initializations). For the 0th- and 1st- order human (left and middle, respectively), all simulations resulted in a single Nash equilibria for each tested human penalty λ_H, indicated by the black bar. For the 2nd-order human (right), simulation outcomes diverged depending on the initialization location, indicated by the bars of varying intensities.

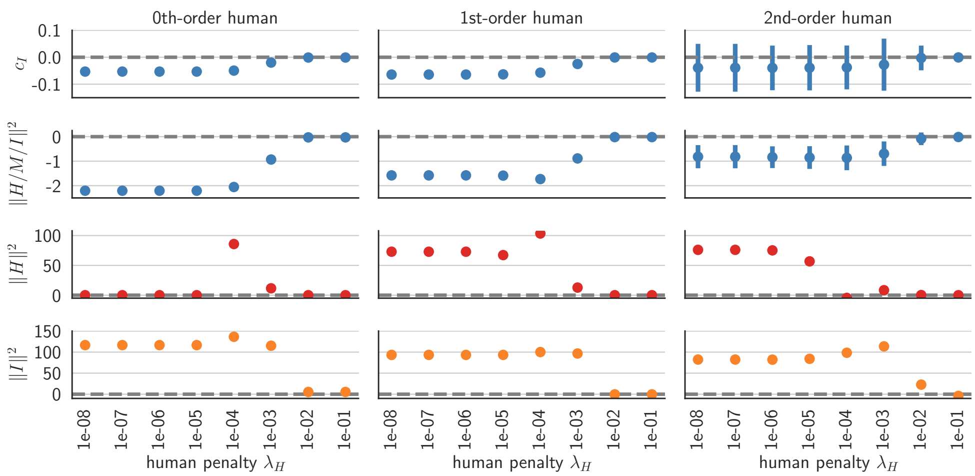

Figure S3: *Simulation results for all interface conditions – performance.* Distributions of differences between performance metrics for baseline and adapted HMI shown using mean +/- standard deviation: (1st row) cost c_I; (2nd row) task performance $\|H/M/I\|^2$; (3rd row) human effort $\|H\|^2$; (4th row) interface effort $\|I\|^2$. The dashed grey line represents the same simulation performance for baseline and adapted HMI. A negative value indicates that the simulated co-adaptation magnitude was lower than the baseline magnitude. A positive value indicates that the simulated co-adaptation magnitude was higher than the baseline magnitude.