human/machine interaction is a *sensorimotor game*

follow along! https://bit.ly/2022-DW

Sam Burden

Assistant Professor

Associate Chair for Diversity, Equity, & Inclusion

Electrical & Computer Engineering

University of Washington

Seattle, WA USA

http://faculty.uw.edu/sburden

human/machine interaction example:
assistive robot

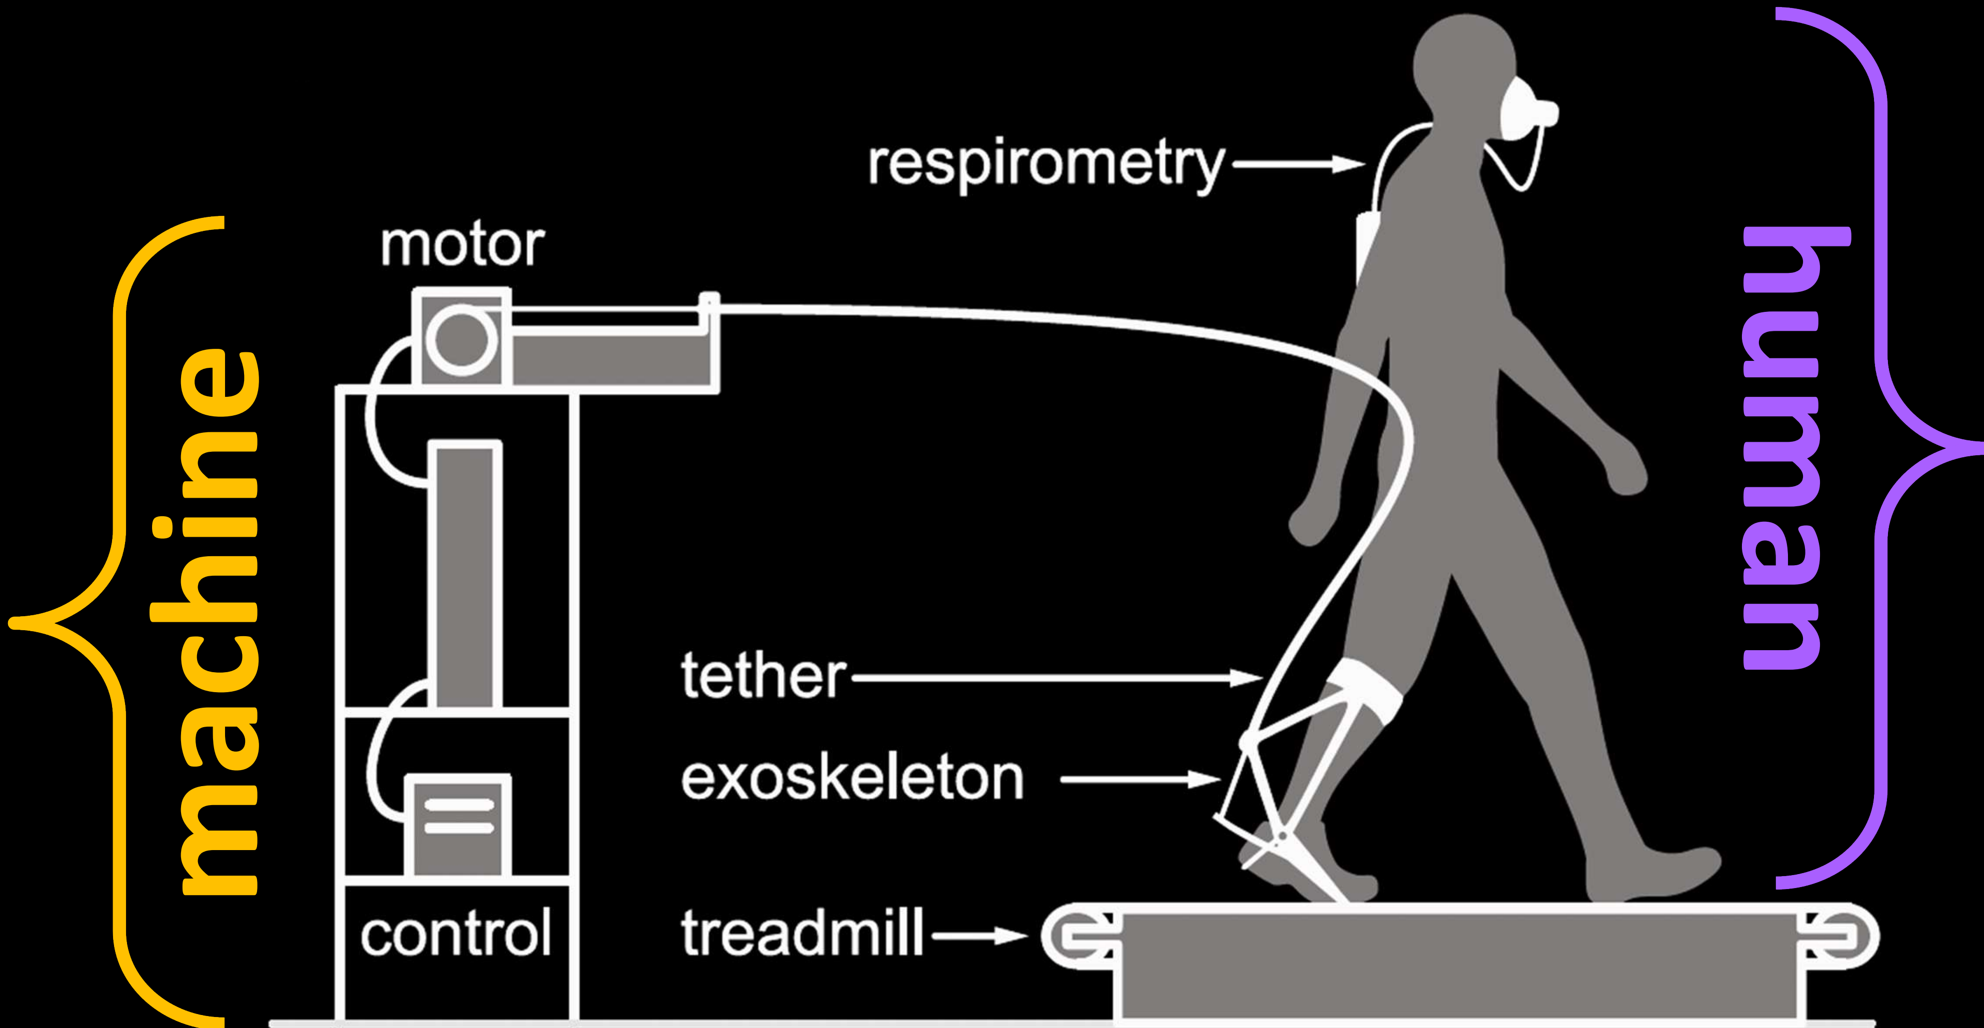

Felt, Selinger, Donelan, Remy, *PLoS One* 2015
Body-in-the-loop: Optimizing device parameters using measures of instantaneous energetic cost
Zhang, Fiers, Witte, Jackson, Poggensee, Atkeson, Collins *Science* 2017
Human-in-the-loop optimization of exoskeleton assistance during walking

human/machine interaction example:
robot teleoperation

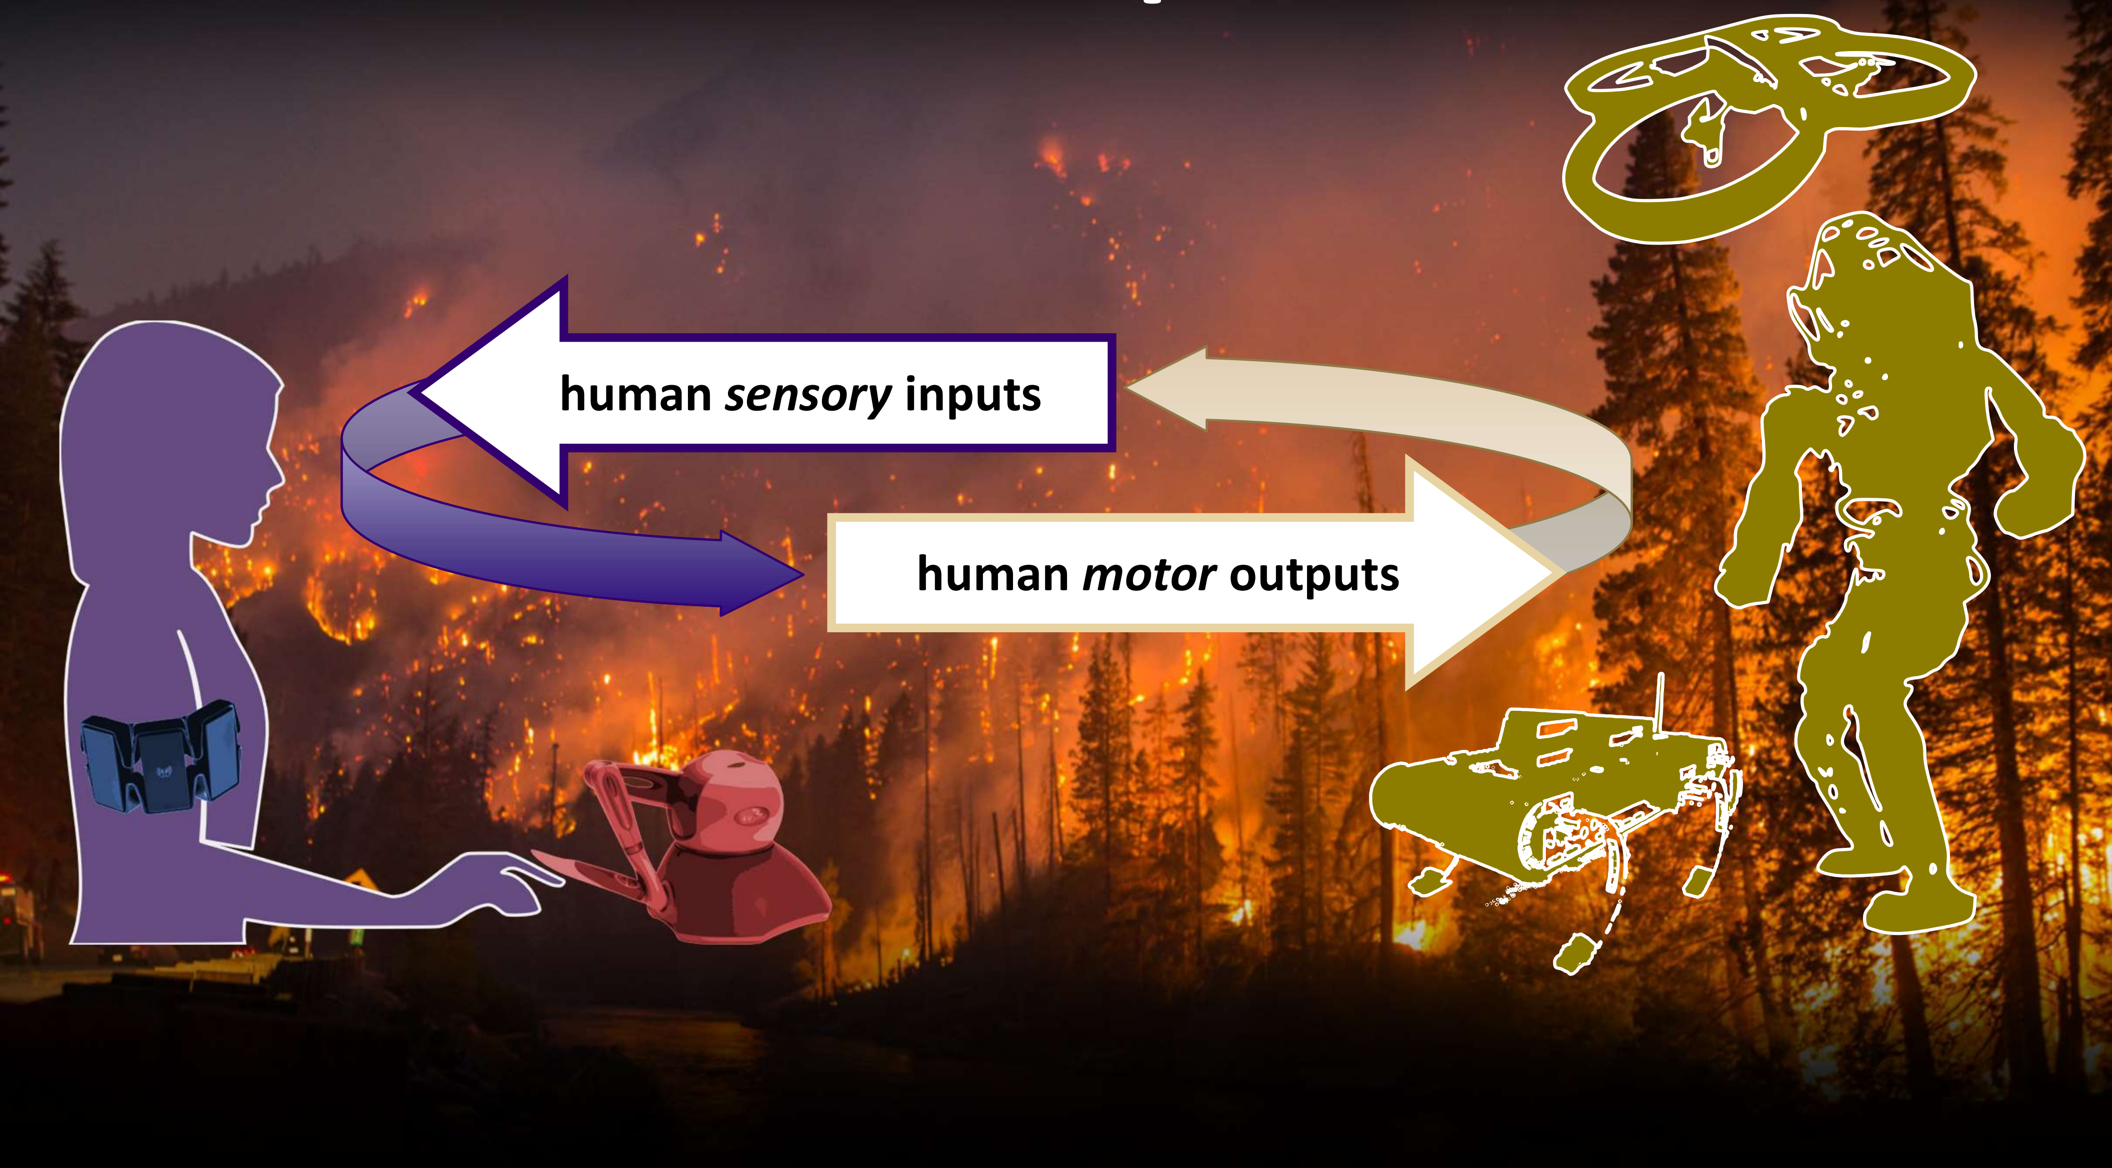

Chiawakum Creek Fire near Lake Wenatchee, WA © Michael Stanford 2015
http://yourshot.nationalgeographic.com/photos/4181903/

human/machine interaction
is a sensorimotor game

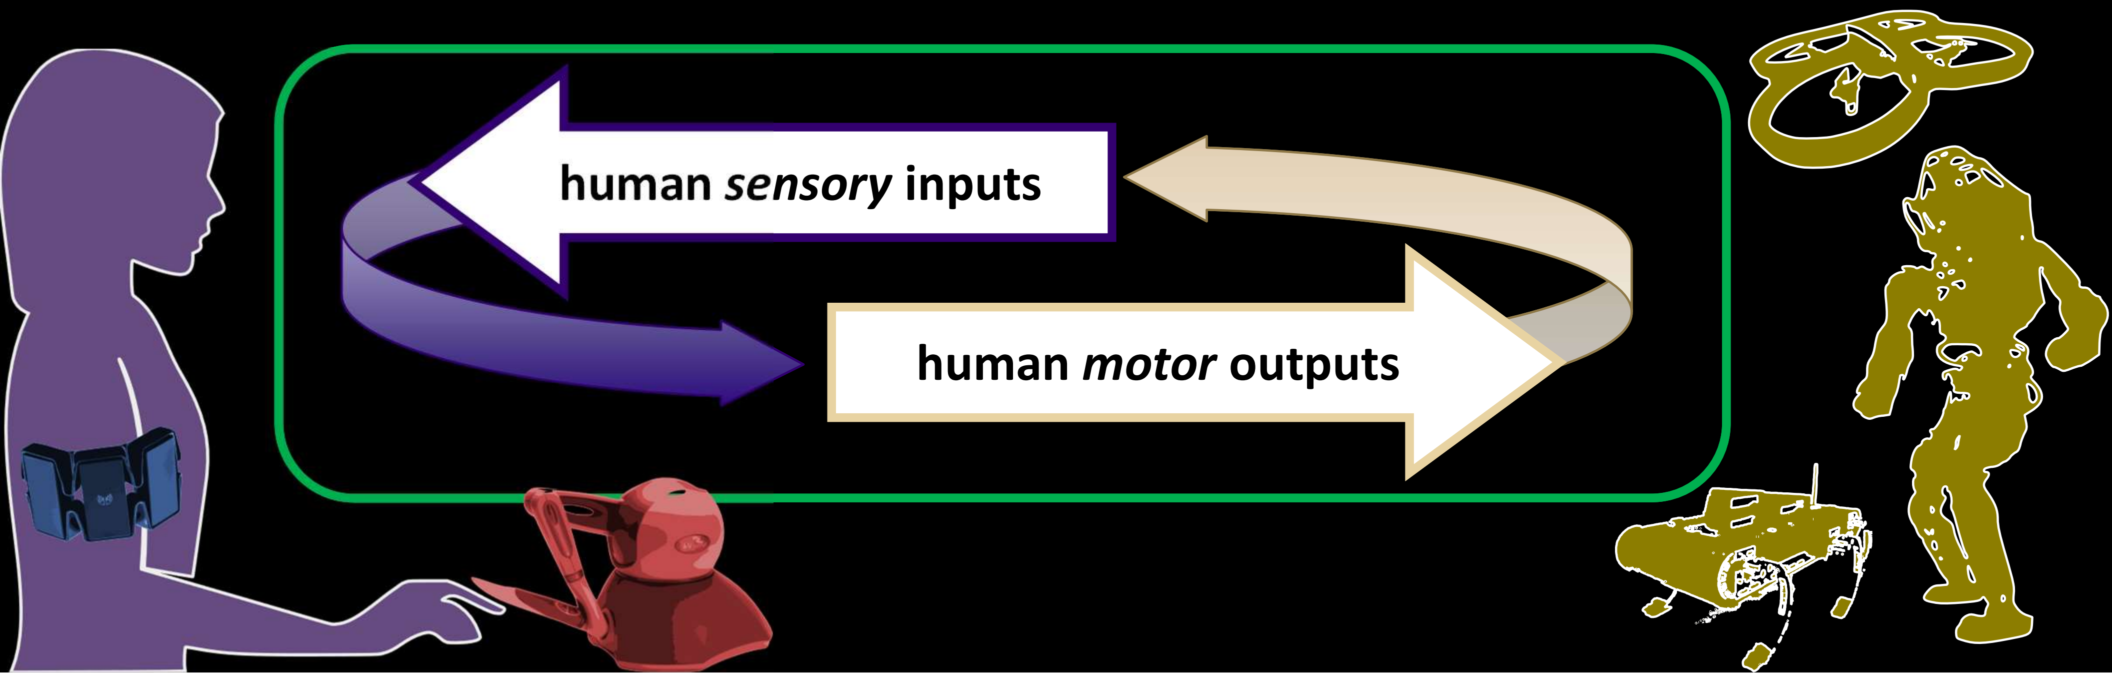

human/machine interaction is a sensorimotor game

Simon *Decision and Organization* 1972
Theories of Bounded Rationality
Russell, Wefald *Artificial Intelligence* 1991
Principles of Metareasoning

*** with "bounded rationality"**

Gershman, Horvitz, Tenenbaum *Science* 2015
Computational Rationality: A Converging Paradigm for Intelligence in Brains, Minds, and Machines
Papadimitriou, Piliouras *ACM SIGecom Exchanges* 2018
Game Dynamics as the Meaning of a Game

observation: humans and machines minimize* interdependent costs

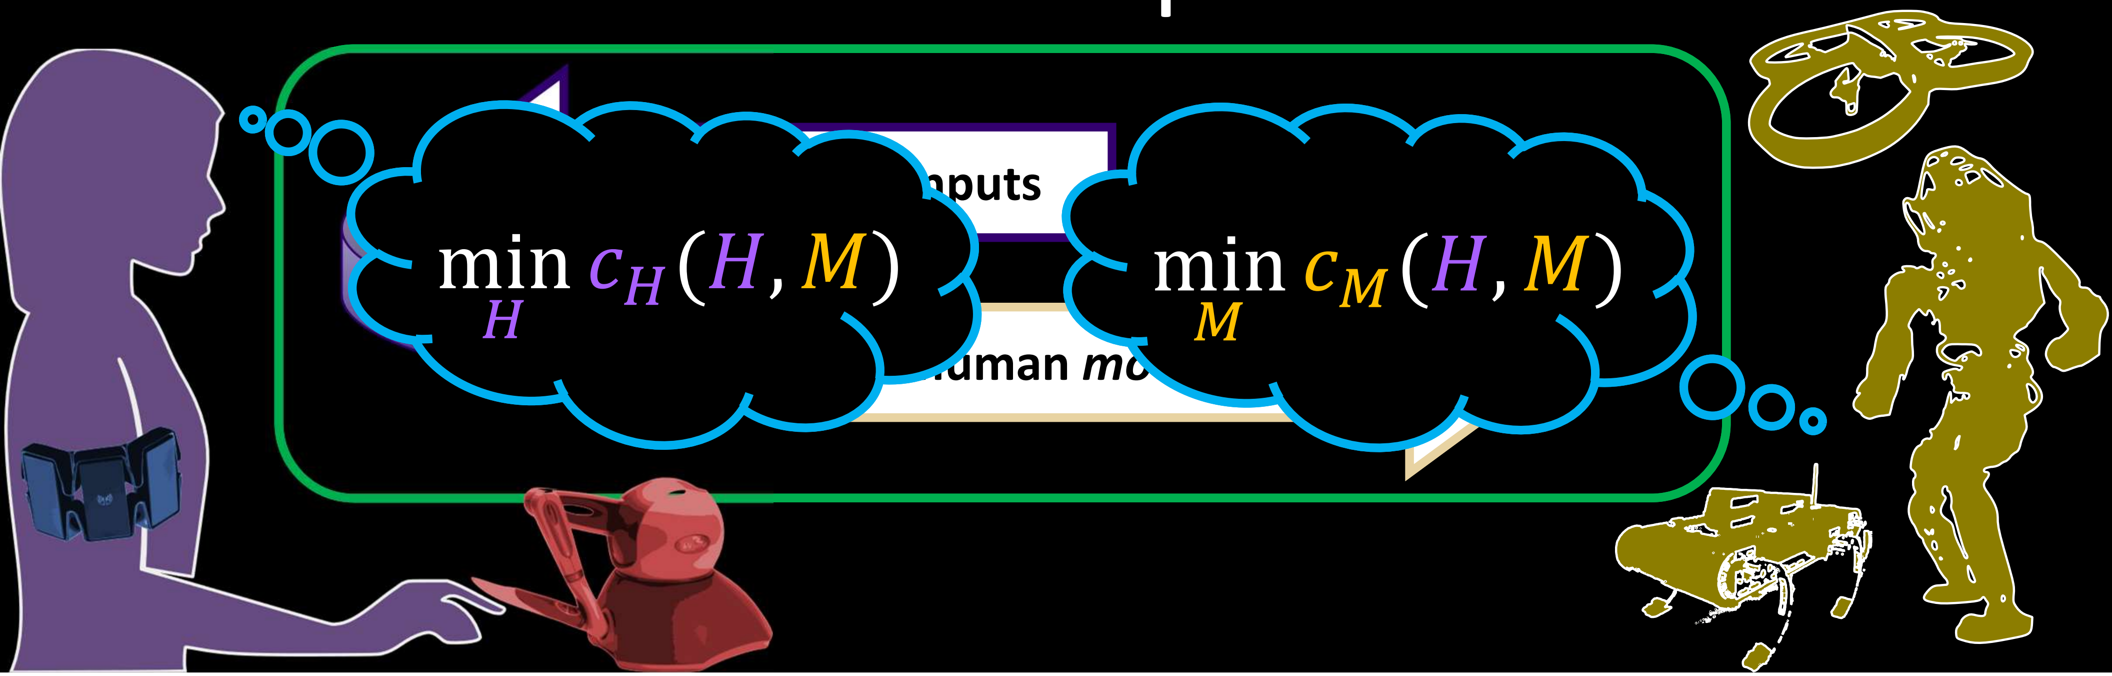

human/machine interaction
is a *sensorimotor game* …

… so what? (why) does it matter ??

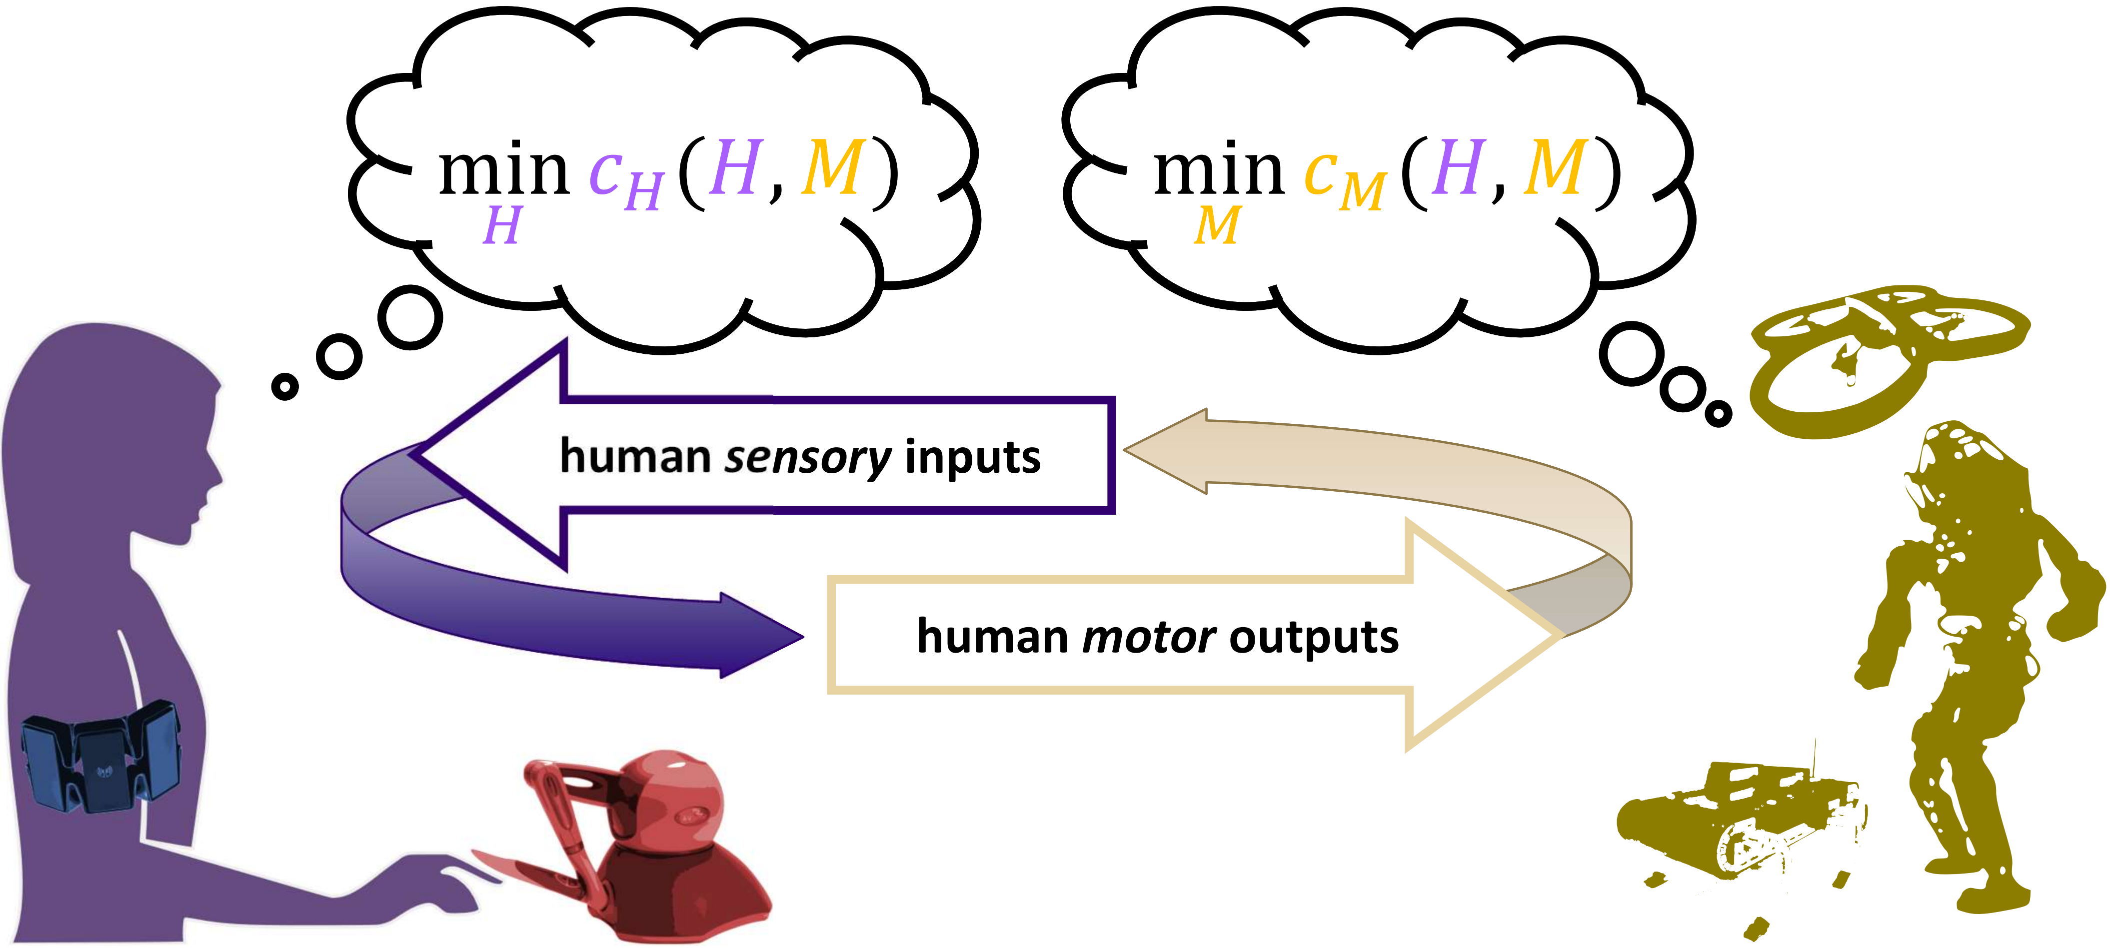

standard algorithms may not work ...

optimization

$$\min_x c(x)$$

standard algorithms like (stochastic) *gradient descent*

$$\dot{x} \approx -\alpha \nabla_x c(x)$$

or *sampling*

$$x^+ \sim P(x; \theta)$$

are **guaranteed to converge** to (local) minimizers of c

Ma, Chen, Jin, Flammarion, Jordan *PNAS* 2019
Sampling can be faster than optimization

game

$$\min_x c_1(x, y) \qquad \min_y c_2(x, y)$$

gradient descent

$$\dot{x} \approx -\alpha_1 \nabla_x c_1(x, y)$$
$$\dot{y} \approx -\alpha_2 \nabla_y c_2(x, y)$$

can easily converge to maximizers, saddles, cycles, or **fail to converge entirely**

Chasnov, Ratliff, Mazumdar, Burden *UAI* 2019
Convergence analysis for gradient-based learning in continuous games

definition of "solution" isn't obvious …

optimization

Bertsekas 1999 *Nonlinear programming*

$$\min_{x} c(x)$$

def: x^* is a *minimum* if deviation increases cost

$$x \neq x^* \Longrightarrow c(x) > c(x^*)$$

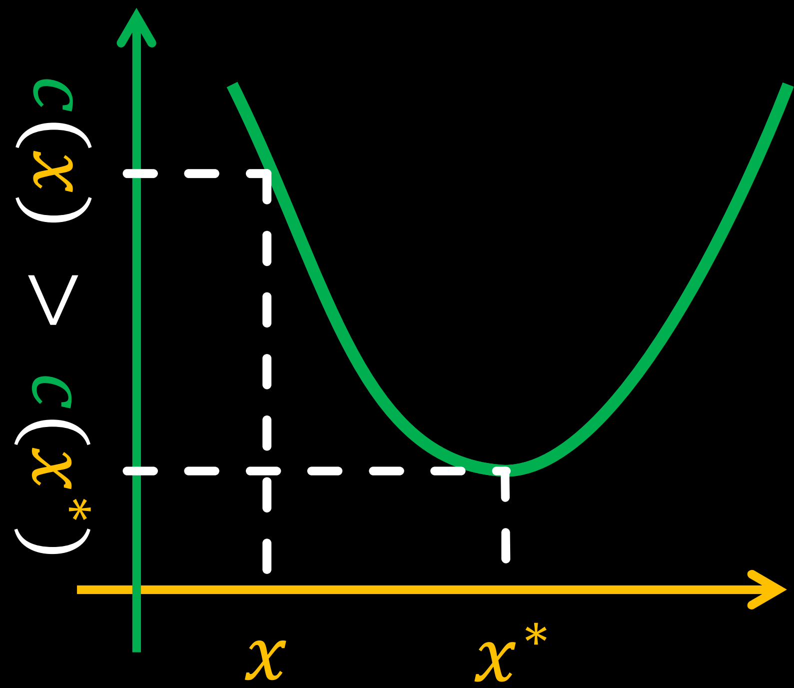

definition of "solution" isn't obvious …

game

Hespanha 2017 *Noncooperative game theory*

— $\min_{x} c_1(x, y)$ — $\min_{y} c_2(x, y)$

● $\min_{x,y} c_1(x, y)$ ● $\min_{x,y} c_2(x, y)$

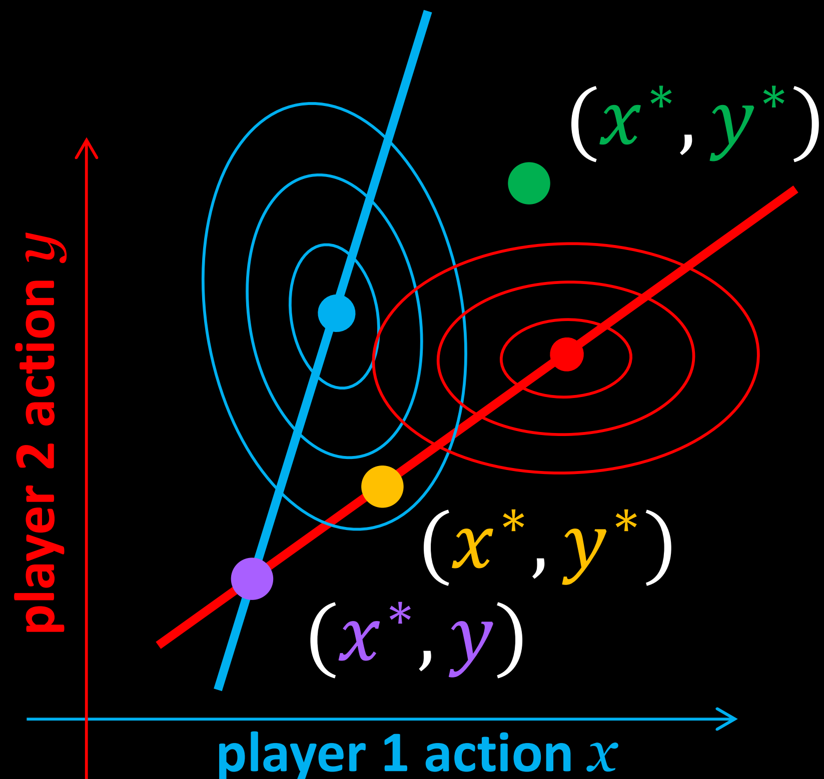

def: (x^*, y^*) is a *Nash equilibrium* if **unilateral** deviation increases cost

$$x \neq x^* \Longrightarrow c_1(x, y^*) > c_1(x^*, y^*)$$

$$y \neq y^* \Longrightarrow c_2(x^*, y) > c_2(x^*, y^*)$$

def: (x^*, y^*) is a (*player 1 led*) *Stackelberg equilibrium* if

$$x^* = \operatorname*{argmin}_{x}\{c_1(x, y^*) \mid y^* = \operatorname*{argmin}_{y} c_2(x, y)\}$$

def: (x^*, y^*) is a *consistent conjectural variations eq.* if …

human/machine interaction is a *sensorimotor game* …

… (how) can machine influence outcomes?

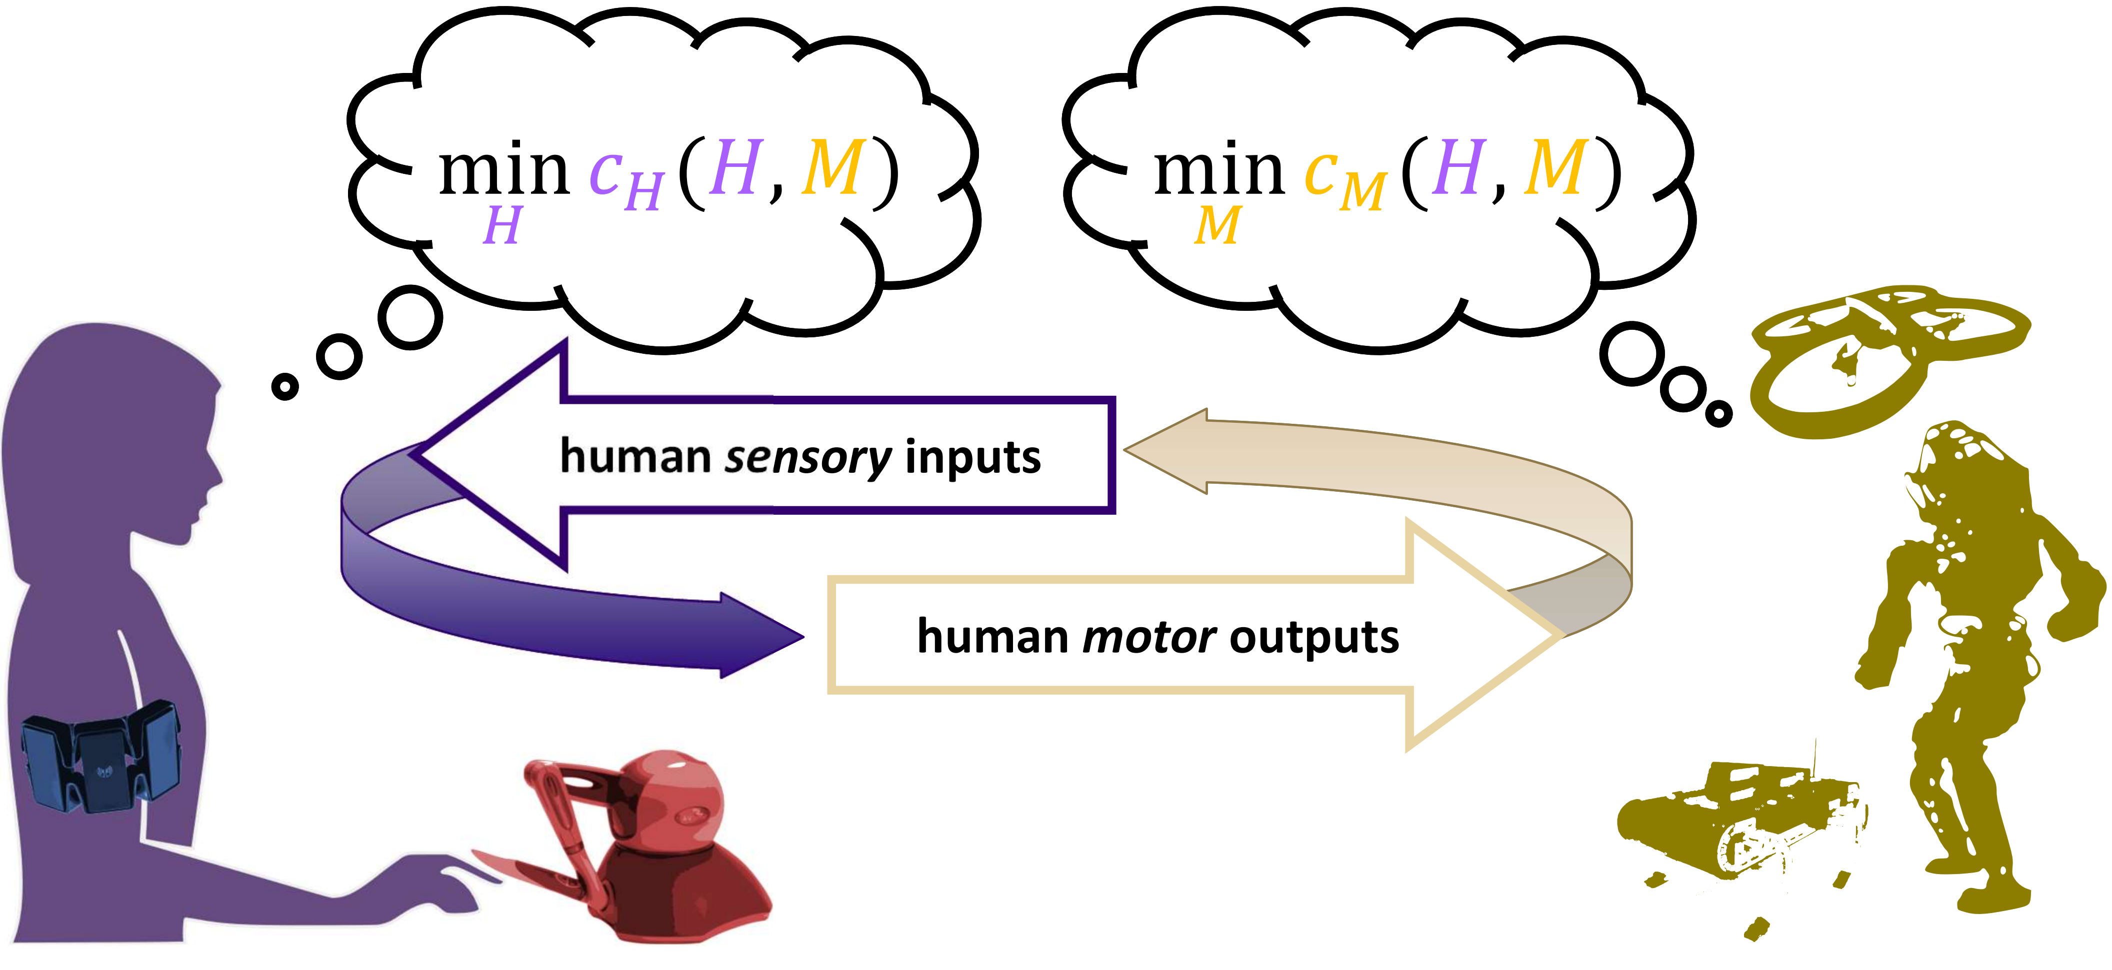

human/machine interaction game

Lillian Ratliff
UW ECE faculty

Ben Chasnov
UW ECE PhD candidate

Ben will be on the academic job market ☺

Prolific

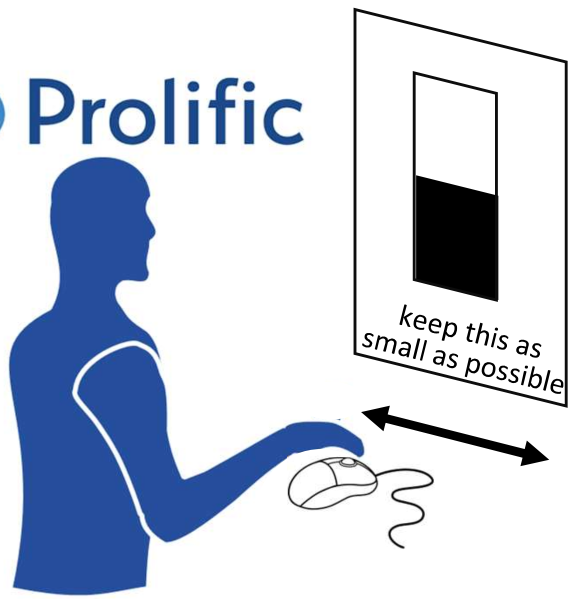

- $N = 20$ participants per experiment
- 10 trials of 40 second duration
- payout ≈ $2 USD

https://dynam.space/study/quadgame.html

experiment 1 (methods):

vary machine's learning rate

human's cost $c_H(x, y)$

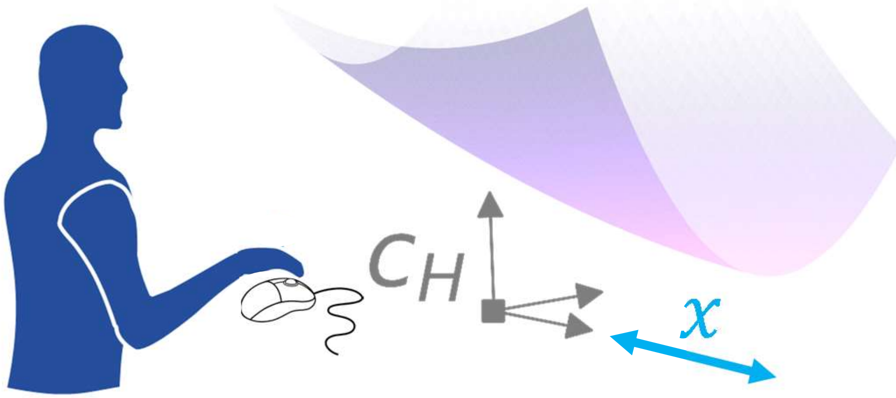

machine's cost $c_M(x, y)$

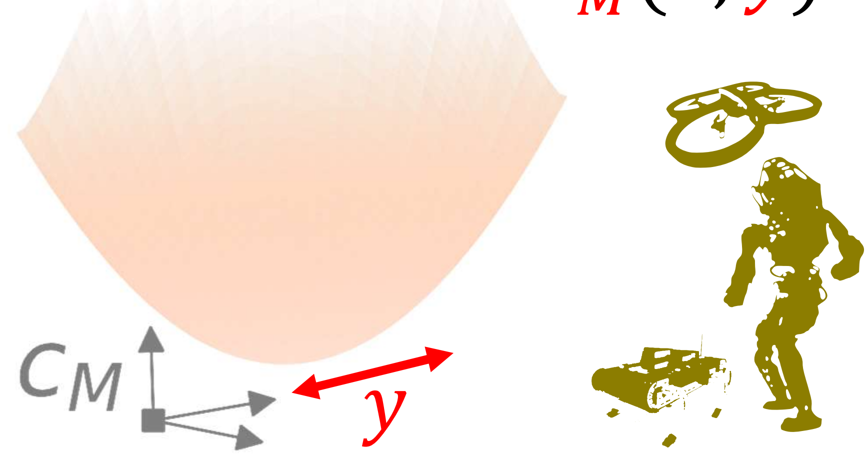

human does ¯_(ツ)_/¯

$$\dot{x} = ???$$

machine does gradient descent

$$\dot{y} = -\alpha \nabla_y c_M(x, y)$$

learning rate changes each trial

$$\alpha \in \{\text{slow}, \text{medium}, \text{fast}\}$$

other details:

- costs are prescribed quadratics and do not change
- machine knows its cost function c_M and human action x
- human only knows $c_H(x, y)$, doesn't know machine action y

experiment 1 (results, $N = 1$):

vary machine's learning rate

human's cost $c_H(x, y)$

human does ¯_(ツ)_/¯

$$\dot{x} = ???$$

machine's cost $c_M(x, y)$

machine does gradient descent

$$\dot{y} = -\alpha \nabla_y c_M(x, y)$$

slow α

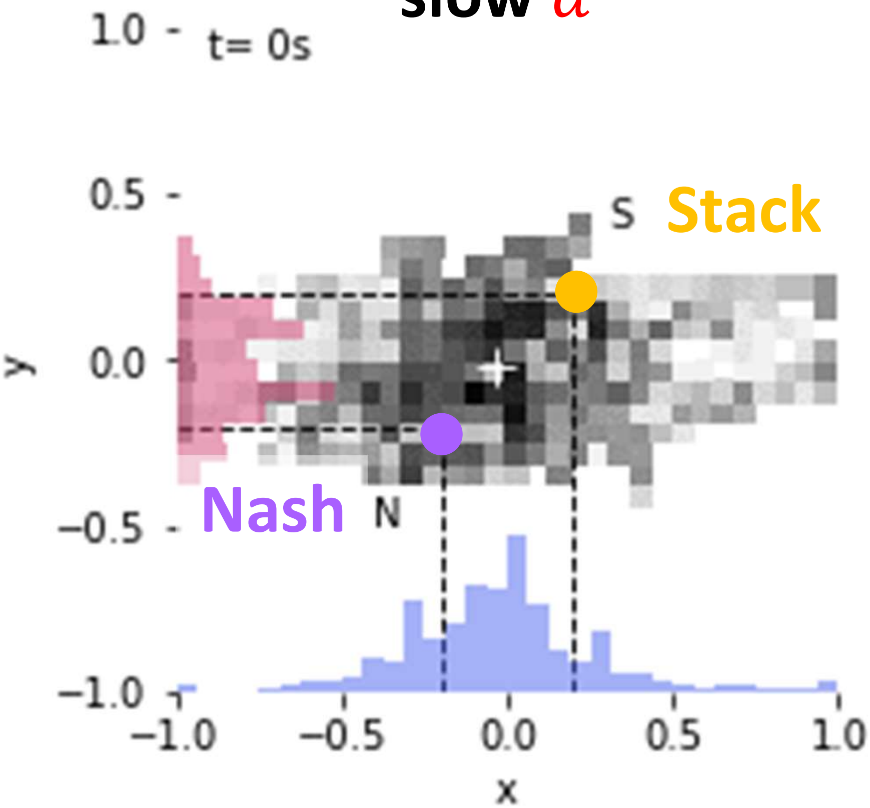

fast α

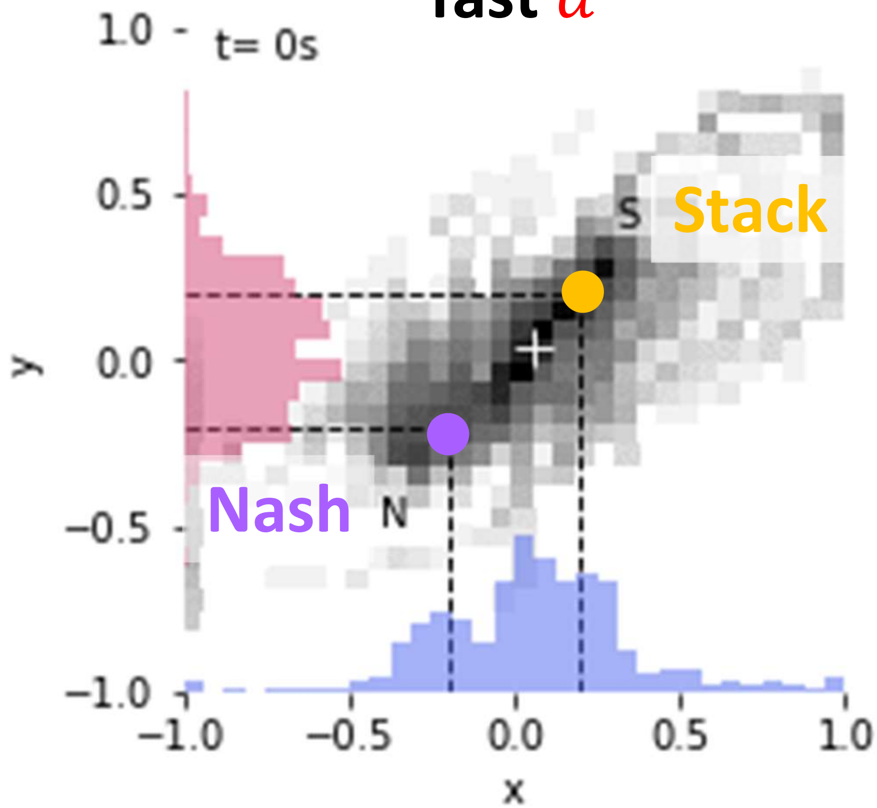


experiment 1 (results, $N = 20$):

vary machine's learning rate

human's cost $c_H(x, y)$

human does ¯_(ツ)_/¯

$$\dot{x} = ???$$

machine's cost $c_M(x, y)$

machine does gradient descent

$$\dot{y} = -\alpha \nabla_y c_M(x, y)$$

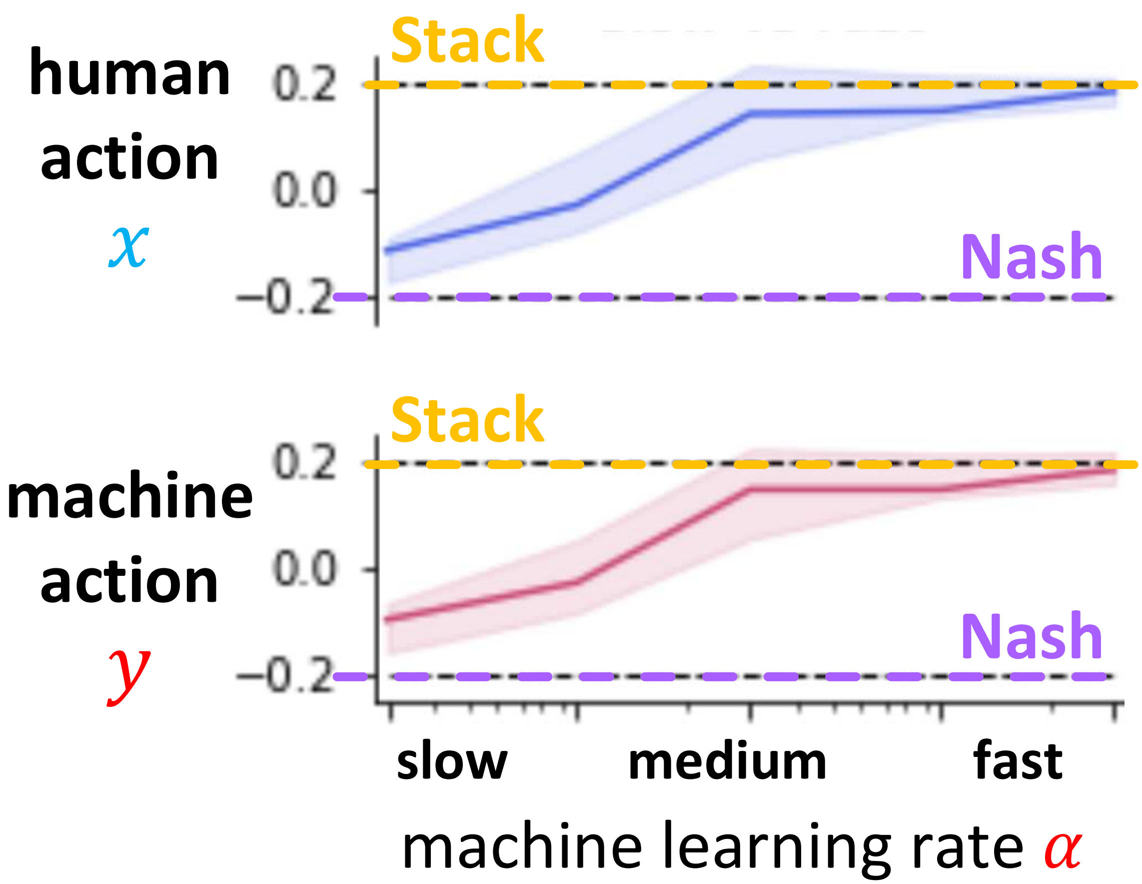

findings:

- increasing machine's learning rate shifts outcome from **Nash** to **Stackelberg** equilibrium

$*$ human **cannot** (only) be doing gradient descent! (learning rates do not change stationary points)

 

experiment 2 (methods):

internal models / *conjectures*

since cost is quadratic, machine's best-response is linear,

$y = L_M(x - x_M^*)$ $\quad x_M^* = x$ coord of M's global min

similarly, natural to hypothesize human responds linearly,

$x \approx L_H(y - y_H^*)$ $\quad y_H^* = y$ coord of H's global min

what if machine estimates L_H?

what if machine uses this
internal model / conjecture
to "outsmart" the human?

$$\min_y \{c_M(x, y) \mid x \approx L_H(y - y_H^*)\}$$

if human responds similarly,
iterating converges to *consistent conjectural variations eq.*

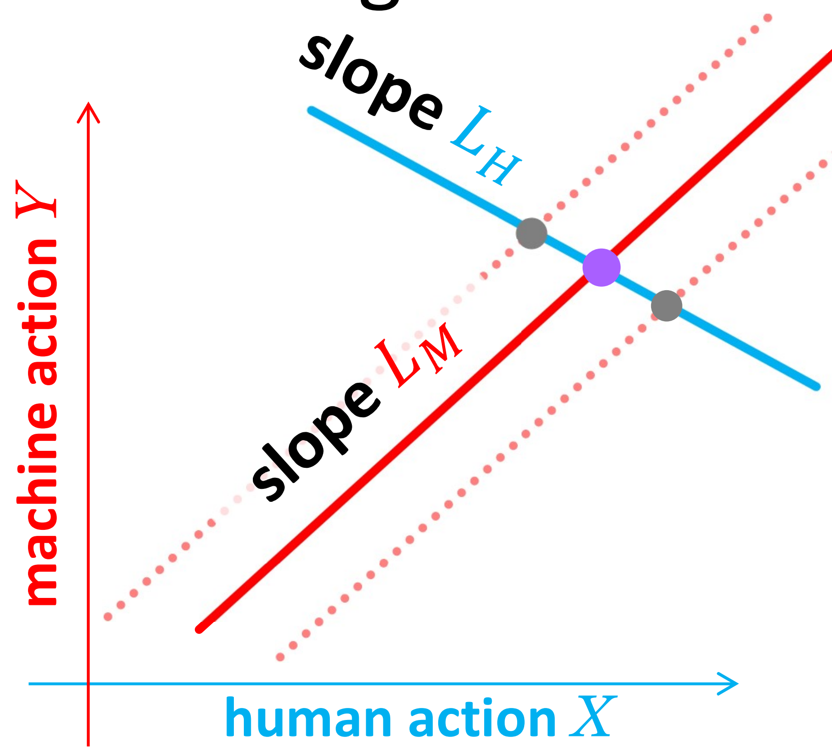


experiment 2 (results, $N = 20$):

internal models / *conjectures*

finding: iterating to "outsmart" …

$$y = L_M^k(x - x_M^*) \qquad x \approx L_H^k(y - y_H^*)$$

$$y = L_M^{k+1}(x - x_M^*) \qquad x \approx L_H^{k+1}(y - y_H^*)$$

$$\dots \qquad \dots$$

$$y = L_M^{\infty}(x - x_M^*) \qquad x \approx L_H^{\infty}(y - y_H^*)$$

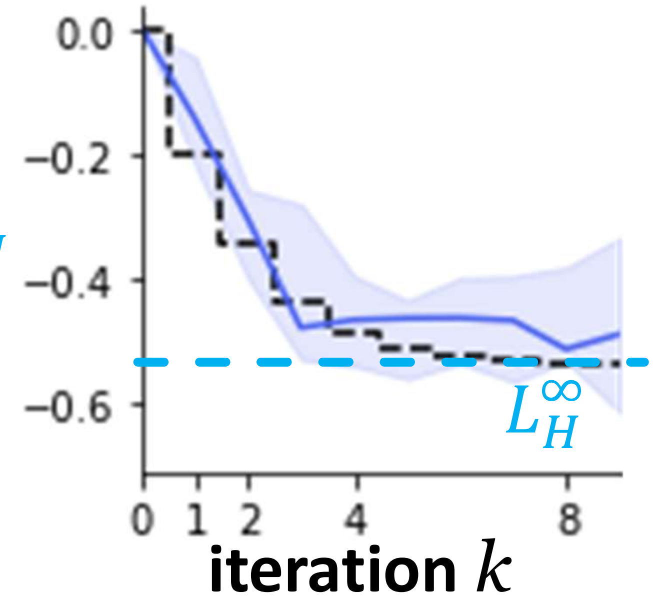

… converges to *consistent conjectures*

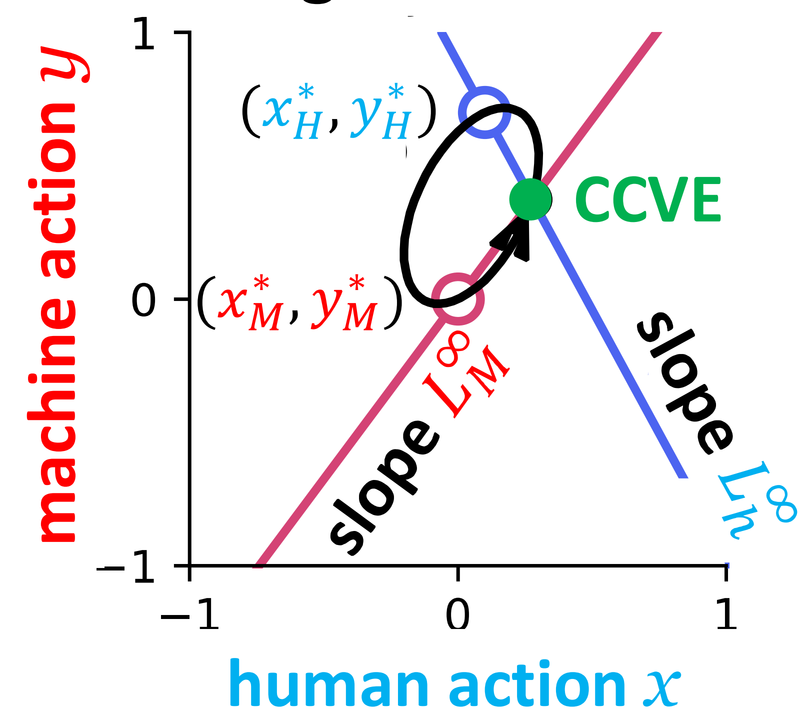

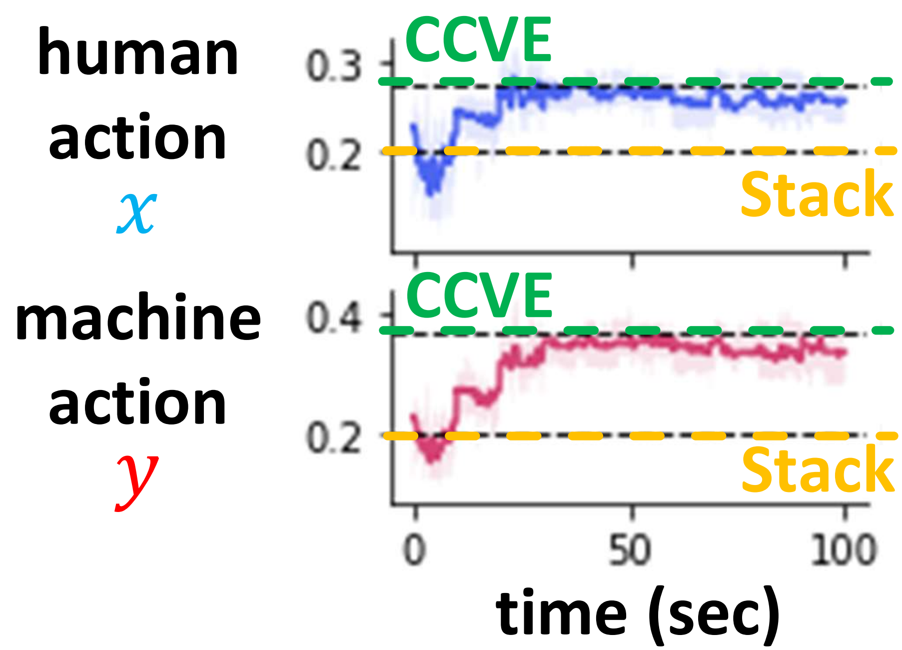


experiment 3 (methods):

machine manipulation

now suppose the machine wants a specific outcome, e.g. its global minimum (x_M^*, y_M^*)

then it can implement a perturbed linear strategy,

$$y = (L_M + \Delta)(x - x_M^*)$$

wait for the interaction to converge (to *reverse Stackelberg* equilibrium),

$$\lim_{t\to\infty}(x(t), y(t)) = (x_\Delta^*, y_\Delta^*)$$

and use this data to descend cost gradient in strategy space,

$$\dot{L}_M = -\alpha \nabla_{L_M} c_M$$

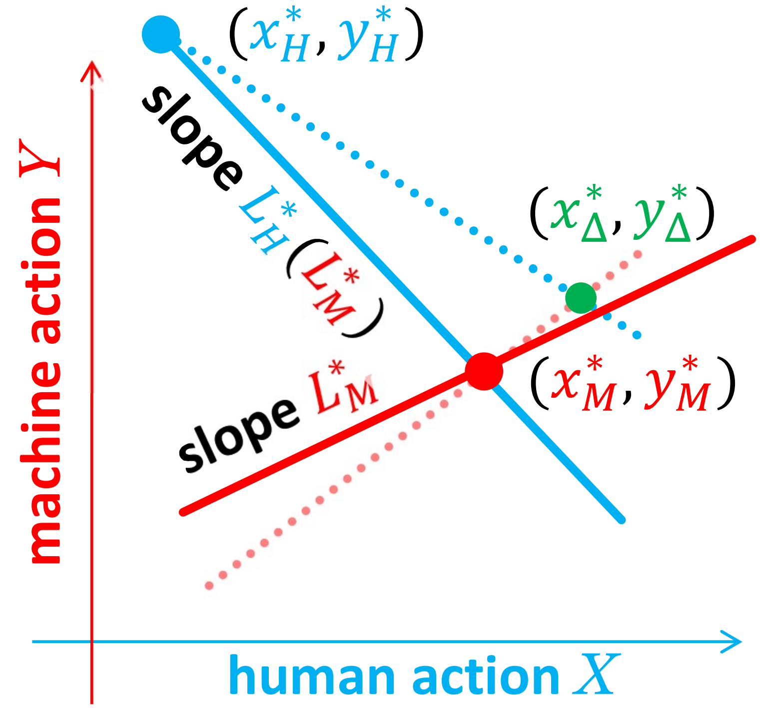

experiment 3 (results, $N = 20$):

machine manipulation

finding:

$*$ machine can "coerce" human to play **any** desired equilibrium using data-driven algorithm,

$$\dot{L}_M = -\alpha \, \nabla_{L_M} c_M$$

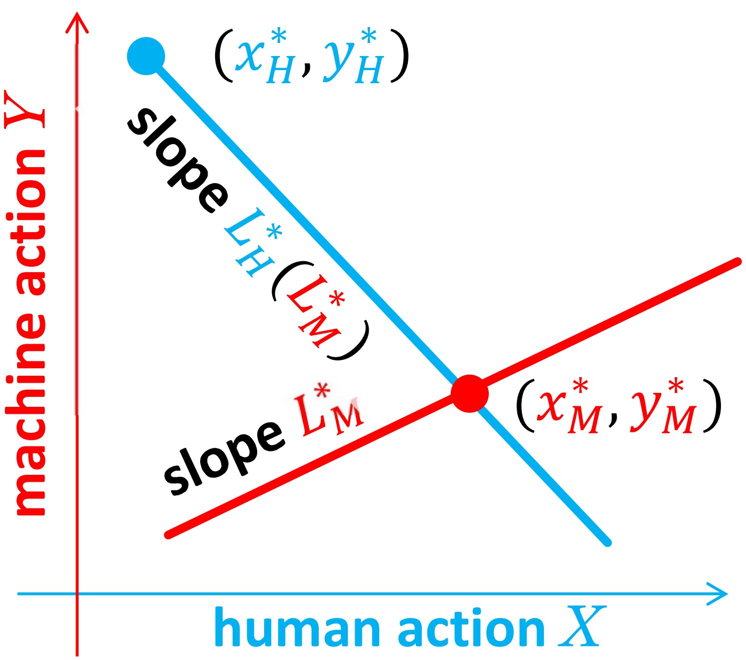

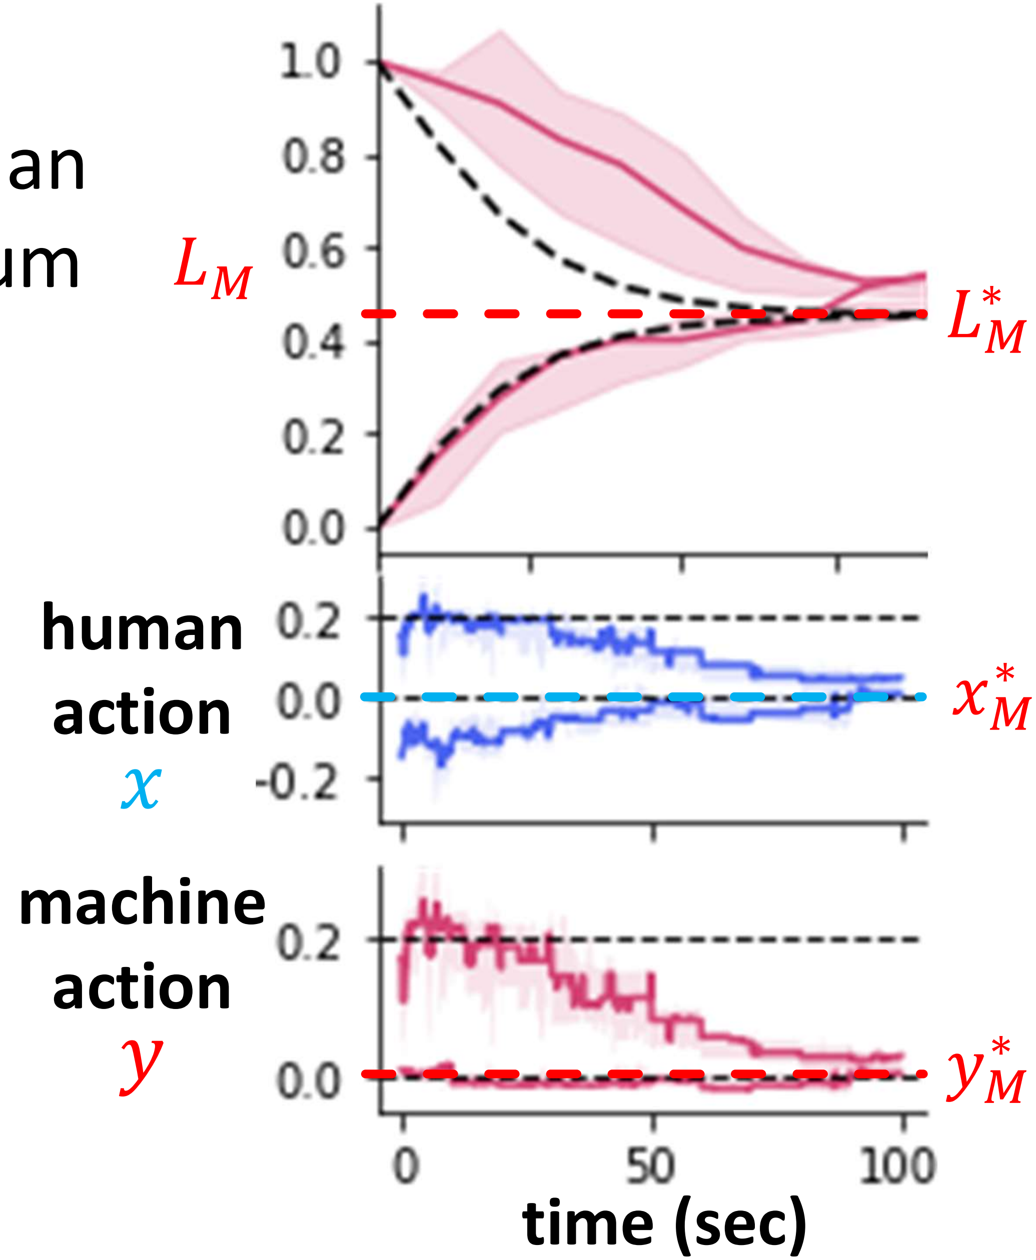

human/machine sensorimotor game:
assistive robot

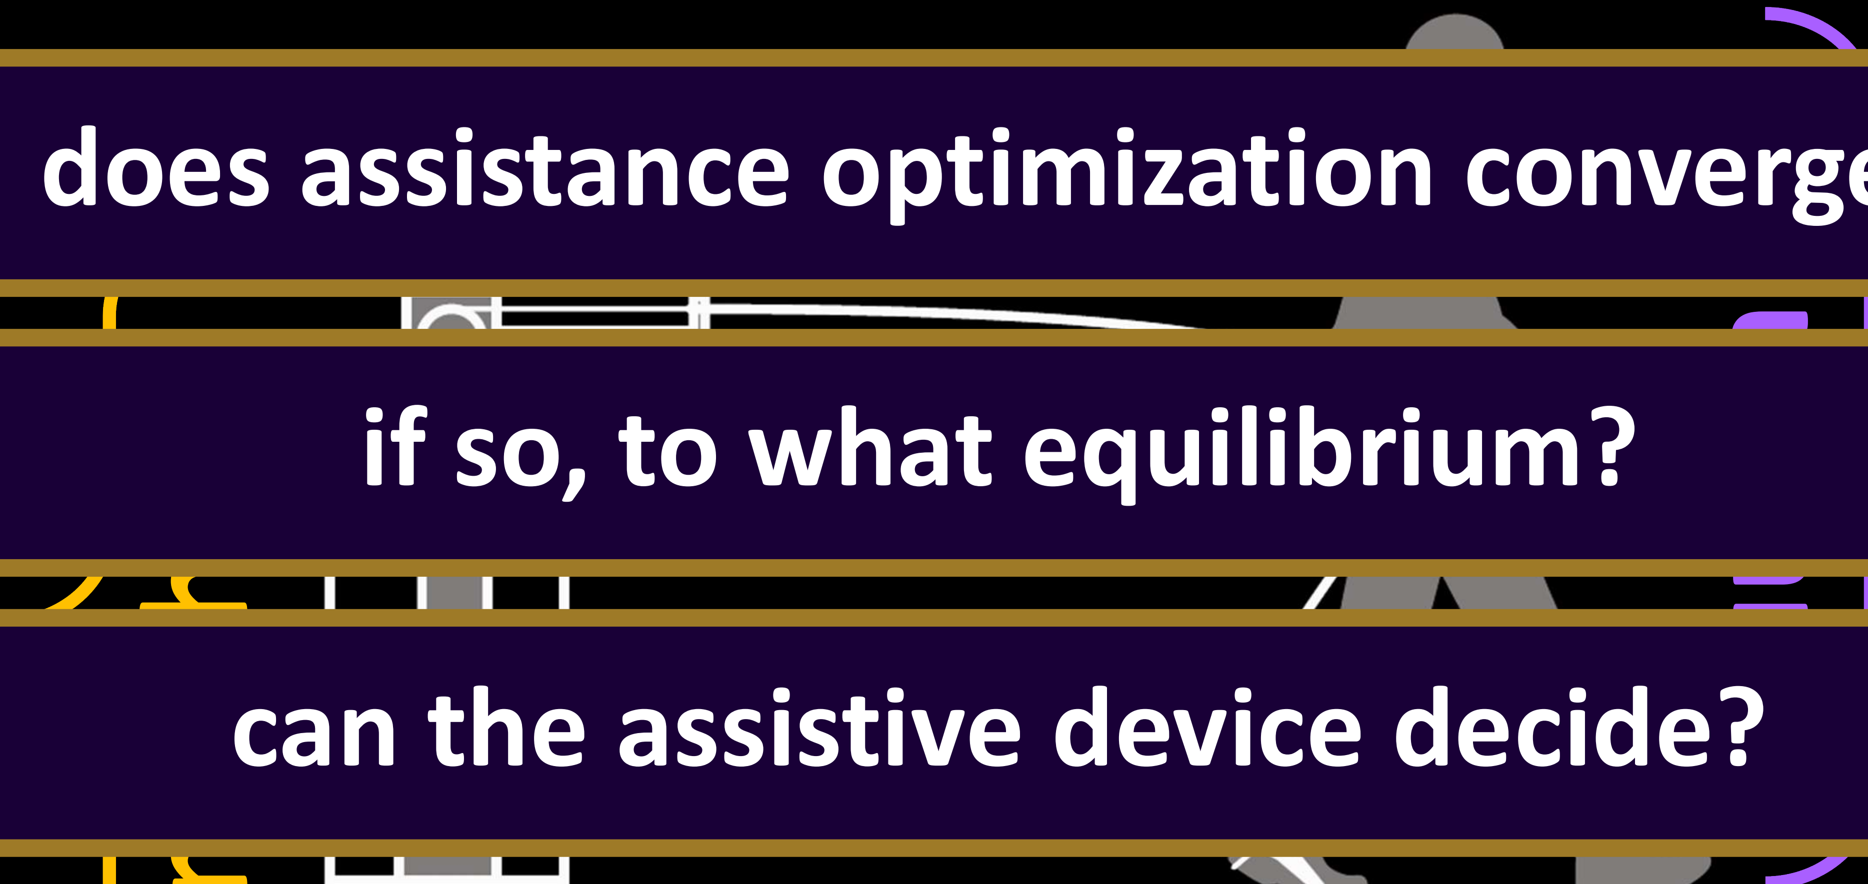

does assistance optimization converge?

if so, to what equilibrium?

can the assistive device decide?

(how) can we converge to H's minimum?

Body-in-the-loop: Optimizing device parameters using measures of instantaneous energetic cost
Zhang, Fiers, Witte, Jackson, Poggensee, Atkeson, Collins *Science* 2017
Human-in-the-loop optimization of exoskeleton assistance during walking

thank you!

sburden@uw.edu
http://faculty.uw.edu/sburden

Lillian Ratliff
UW ECE faculty

Ben Chasnov
UW ECE PhD candidate

NSF M3X CAREER #2045014:
Human/Machine Collaborative Learning and Control of Contact-Rich Dynamics

NSF CPS Medium #1836819:
Certifiable reinforcement learning for cyber-physical systems

findings:

- when H and M play a game
- the outcome is never the same
  - there are so many ways
  - to outsmart with plays
- that I can't recall all their names …